From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Economics

2019

Essays in Labor Economics and Innovation

Cramer, Judd Nathan Levine

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp01w37639557

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Essays in Labor Economics and Innovation

Judd N. L. Cramer

A Dissertation Presented to the Faculty of
Princeton University in Candidacy
for the Degree of Doctor of Philosophy

Recommended for Acceptance
by the Department of Economics
Advisors: Henry S. Farber and Alan B. Krueger

January 2019

© Copyright by Judd N. L. Cramer, 2019.

All rights reserved.

Abstract

This collection of essays applies empirical techniques to questions in labor economics, with a focus on innovation. A common theme among these articles is that changes in technology and institutions can have impacts on the functioning of the labor market and innovation.

Chapter 1 explores an understudied question in innovation: what is the relative importance of patent term (i.e., duration of patent protection) across different industries? This article[1] measures term sensitivity through patent applicants' response to the passage of the TRIPS agreement, which changed how term was calculated. We find significant differences in term sensitivity across industries, some of which follows conventional wisdom (patent term is important in pharmaceuticals) and some of which does not (it also matters for software).

Chapter 2 investigates whether the emergence of ride-sharing companies like Uber and Lyft lowered the wages and employment of taxi drivers and chauffeurs. To evaluate the impact of ride-sharing services, this paper[2] uses the timing of the rollout of Uber along with the company's degree of popularity across cities to estimate the effects of competition from Uber on taxi and limousine drivers' employment and wages. In contrast to the literature on disruptions in other industries, including the deregulation of the trucking and airline industries in the 1970s, I find no evidence that the advent of Uber has had a negative impact on wages in the taxi and limousine industry thus far.

[1] *Who Cares About Patent Term? Cross-Industry Differences in Term Sensitivity* (co-authored with fellow Princeton Ph.D., Neel Sukhatme).

[2] *The Effect of Uber on Employment and Wages of Taxi and Limousine Drivers.*

Chapter 3[3] documents that it was the behavior of the record number of long-term unemployed during and after the Great Recession that explained the initial movements away from the Beveridge curve, which puzzled economists. I extend the calibrated model originally estimated in Krueger, Cramer, and Cho (2014) and focus on the channels through which duration dependence could be a factor. Incorporating the cyclical patterns in the labor force flows of the long-term unemployed (lower labor force withdrawal rates during recessions and higher labor force withdrawal rates as recoveries continue) improved the accuracy of the model's forecast of the unemployment rate and produced the counterclockwise loops as predicted in Blanchard and Diamond (1989).

[3] *Long-term Unemployment and the Loop in the Beveridge Curve: The Importance of Cyclical Labor Force Withdrawal by Duration.*

Acknowledgments

Obtaining a Ph.D. is a long and arduous process. In my case, I made the process a good deal longer than necessary. Without the patience and support of so many people, this journey would never have been possible. First, I am deeply grateful to my advisors, Hank Farber and Alan Krueger. Alan and Hank have supported me throughout my time at Princeton, and their mentorship has been paramount to my success. I hope to one day possess a fraction of their knowledge and assist the next generation of researchers. I am also grateful to the other member of my dissertation committee Orley Ashenfelter, a giant of Labor Economics, who was always willing to talk, and was my favorite teacher in the program.

In addition, I thank the rest of the faculty, students, and staff at Princeton's Industrial Relations Section, where my office was located. I am grateful for financial support from the Clarence J. Hicks Memorial Fellowship in Industrial Relations. More importantly, I appreciate all the feedback I've received over the years, both at lunch seminars and in more informal discussions. I am also thankful for the continuous support of the IR Section staff, particularly Linda Belfield, Patti Tracey, and Jeannie Moore. I have caused them many headaches over the years, but they have been nothing but courteous, efficient, and forgiving. Laura Hedden is an incredible resource for all graduate students in the ecomomics department and helped me immensely throughout this process.

My classmates at Princeton were the best peers of whom I could have dreamed and their contributions to my development as an academic and person were immense. I would not have survived first year without Cheng Chen. Brice and Benjamin were two of the smartest people I have ever met. I think of Tom Carter and cannot help but smile. My officemates over the years: Pauline Leung, Inessa Liskovich, and Neel Sukhatme provided much assistance and

put up with my sloppiness. To my group of David Cho, Scott Kostyshak, Andres Maggi, Neel Sukhatme, and Tom Winberry, in particular, I owe a debt that could not be paid in 10,000 years, regardless of interest rate. I also made other great friends in Princeton. I thank Robin Gomila, whose optimism about his own research agenda rubbed off on me. Jessica Wagner has been instrumental in the preparation of this document as well as my ability to persevere through this process in general. Her support and love have allowed me to embrace struggle and failure more than I thought I could stomach. I hope I can pay her back.

I took repeated hiatuses from Princeton to work as a staff economist for the Council of Economic Advisors (CEA) at the White House from 2011 forward. It was the honor of a lifetime to serve under President Obama and wonderful Chairs: Austan Goolsbee, Alan Krueger, and Jason Furman. I am thankful that two other professors with CEA connections, Harvey Rosen and Owen Zidar, agreed to be the non-readers on my committee. Harvey was a former chair of the CEA and Owen walked me through his legendary spreadsheets on indebetedness my first day on the job in Washington. All of my junior colleagues for the year at the CEA were indefatigable and uber intelligent. If I am ever approaching a fiscal cliff, my team is Matt Aks, Jeff Borowitz, Colleen Carey, David Cho, Sandra Levy, Carter Mundell, Pedro Spivakovsky-Gonzalez, Reid Blake Stevens, Seth Werfel, and Julia Yoo every single time.

I am incredibly fortunate to have left the dream teams at Princeton and the CEA for the A-team at Harvard. Teaching and advising alongside Thomas Baranga, Greg Bruich, Anne Le Brun, Dani Doyle, Kiran Gajwani, Annie Gajwannie, and Jonas Poulsen has been a wonderful experience. If I had not wound up with such generous, smart, and encouraging colleagues, I do not know where I would be. Jeffrey Miron has showed me extreme patience and I am happy to endorse him for governor at a date to be announced.

Finally, my family has been a continual source of support and inspiration for me. I am indebted to my late paternal grandparents, Sylvia and Milton Cramer. I am incredibly

lucky to have such a close relationship with my maternal grandparents, Jackie and Howard Levine. Thinking of the tremendous sacrifices that all my forebearers made to allow me to grow up in previously unimaginable affluence reminded me how silly it was to be intimidated by running robustness checks. I am immensely honored and humbled to be the first to receive a doctorate on their behalf. I am so lucky to be very close with the rest of my family and thank all my aunts, uncles, and cousins for their support. I think of my late Great (in many ways) Uncle Steve Solarz often whose devotion to higher learning set me on the path I currently travel. My mother and sister have dealt with me for many years and have always shown me unconditional love. I have always been encouraged to pursue my dreams and have never once worried what others would think. It is the regret of my life that I was unable to know my father as an adult. The stories I have heard about his humor, intelligence, and goodness, however, have always served as a cynosure in my life. It is to his memory this work is dedicated. There are far more people than I could ever begin to write about on this page. I hope I will have the opportunity to tell them in person. Thank you.

In memory of Andrew Ira Cramer.

Contents

Abstract iii

Acknowledgments v

List of Tables xiii

List of Figures xvi

1 Who Cares About Patent Term? Cross-Industry Differences in Term Sensitivity . 1

1.1 Introduction 1

1.2 Patent Term and Innovation 4

1.3 Measuring Patent Applicant Delay 9

1.3.1 Dataset and Empirical Strategy 10

1.3.2 Summary Statistics and Non-Parametric Graphs 14

1.3.3 Results Using Applicant Delay Measure 18

1.3.4 Results Using Disaggregated Measures of Patent Delay 26

1.3.5 Cross-Industry Differences in Prosecution Delay 30

1.3.6 Cross-Check Using Patent Count Measure 35

1.3.7 Other Changes in TRIPS 39

1.4 Model of Patent Prosecution Time 41

1.4.1 Model with Fixed Patent Term 44

1.4.2 Model with Term Dependent on Prosecution Time 45

1.4.3 Differences in Prosecution Time Under the Two Patent Term Models 46

1.5 Maintenance Fees as a Measure of Late-Term Patent Value 50

1.6 Conclusion 56

A Additional Figures and Tables 58

B Proofs 62

B.1 Proof of Lemma 1 62

B.2 Proof of Proposition 1 62

B.3 Proof of Proposition 2 64

References 65

2 The Effect of Uber on the Employment and Wages of Taxi and Limousine Drivers 68

2.1 Introduction 68

2.2 Background 71

2.2.1 Background on Uber 71

2.2.2 Background on the Taxi and Limousine Industry 74

2.3 Related Literature 76

2.4 Model 80

2.4.1 Competitive market with unilateral pricing of lease rates, c in the market 81

2.4.2 Rent-sharing case, with Nash bargaining over profits 81

2.5 Data Sources 83

2.6 Empirical Work 86

2.7 Conclusion 91

References 93

3 Long-term Unemployment and the Loop in the Beveridge Curve: The Importance of Cyclical Labor Force Withdrawal by Duration 110

3.1 Introduction 110

3.2 Update on Those on the Margins of the Labor Force 113

3.2.1 Labor Force Performance of the Unemployed 113

3.2.2 Labor Force Performance of Nonparticipants 116

3.3 The Beveridge Curve 117

3.3.1 Previous Models 118

3.3.2 Model Extension 121

3.3.3 Simulation and Updates 123

3.4 Summary and Conclusion 127

References 129

List of Tables

1.1 Effect of TRIPS on Applicant Delay 22

1.2 Effect of TRIPS on Applicant Delay with Time Trend Corrections 24

1.3 Effect of TRIPS on Applicant Requested Extensions 28

1.4 Effect of TRIPS on Non-Final Office Action Response Time and Issue Fee Payment Time 30

1.5 Difference-in-Difference With Count Ratio Measure 38

1.6 Post-TRIPS Coefficient and % Maintained Across Subcategories: Chemicals, Computers & Communications, and Drugs & Medical 53

1.7 Post-TRIPS Coefficient and % Maintained Across Subcategories: Mechanical, Electrical, and Others 54

1.8 Difference-in-Difference with Maintenance Fee Measure 55

1A.1 Applicant Delay by Category with Time Trends 59

1A.2 Patent Classes w/Post-TRIPS Coeff. Significant at 5% Level 60

1A.3 Patent Classes w/Post-TRIPS Coeff. Significant at 5% Level (cont'd) 61

2.1 Uber effect on wages and employment in Uber's 18 largest markets (2008 to 2014) 99

2.2 Relationship between Uber drivers and Uber-related Google search activity . 99

2.3 Two-sample Two Stage Least Squares estimates of popularity of Uber on taxi and limousine wages and employment 100

2.4 Difference-in-Difference estimates of the effect of the length of Uber's presence on employment and earnings of limousine and taxi drivers 101

2.5 Relative time models for Uber 102

2.6 The effect of Uber on taxi and limousine wages and employment using truck drivers as a control group 103

2.7 The effect of Uber on taxi and limousine wages and employment using all other employment as a control group 104

3.1 Forecasts of the Unemployment Rate and Actual Rate For All Workers 16+, 2014-2017 134

List of Figures

1.1 Effect of TRIPS on Patent Term 9

1.2 Sample Transaction History 11

1.3 Pendency and Applicant Delay for Applications Filed in 1994-1996 15

1.4 Pendency by Category 18

1.5 Applicant Delay by Category 19

1.6 Local Polynomial Test for Trend Break in Applicant Delay 20

1.7 Moving Cutpoint Test with Standard Error Bands 26

1.8 Cross-Industry Differences in Applicant Delay (Post-TRIPS coeff.) 31

1.9 Moving Cutpoint Test Across Industry Categories 34

1.10 Number of Applications Filed by Category 36

1.11 Example of Patent Profits Evolution Over Time 43

1.12 Example of Prosecution Cost and Prosecution Time Relationship 43

1A.1 Relationship Between Issue Fee Payment Date and Average Patent Issuance Date 58

2.1 Drivers on the Uber System in the US and Google Searches for "Uber", Jun. 2012 to Oct. 2015 105

2.2 Active Uber Drivers by City, Jun. 2012- Dec. 2014 (Hall and Krueger, 2015) 106

2.3 Uber Drivers vs. Uber Search Intensity 107

2.4 Uber Drivers vs. Uber-related Google Searches 108

2.5 Total Taxi Employment, Jan. 2008-Jun. 2017 (Current Population Survey) . 108

2.6 Total Taxi and Limousine Hours Worked, Jan. 2008-Jun. 2017 (Current Population Survey) 109

3.1 (a) Job Openings vs. Unemployment Rate, Dec-00 to Apr-14 135

3.1 (b) Job Openings vs. Short-Term Unemployment Rate, Dec-00 to Apr-14 . . 136

3.2 Job Openings vs. Unemployment Rate, Dec-00 to Apr-18 137

3.3 Of Total Unemployed, Share Unemployed over 26 Weeks 138

3.4 Monthly Probability of Transitioning from Unemployment to Employment by Duration of Unemployment 139

3.5 Monthly Probability of Transitioning from Unemployment to Employment by Duration of Unemployment 140

3.6 Monthly Probability of Transitioning from Unemployment to Out of the Labor Force by Duration of Unemployment 141

3.7 Monthly Probability of Transitioning from Unemployment to Out of the Labor Force by Duration of Unemployment 142

3.8 Transition Rate from Not in Labor Force 143

3.9 Calibrated Beveridge Curve, 2008-2013 144

3.10 Calibrated Beveridge Curve with 2006 to 2007 Average Labor Force Exit Rates 145

3.11 KCC Beveridge Curve Projection and Actual Data 146

3.12 KCC Beveridge Curve Projection and Realized Vacancy Rate and Actual Data147

Chapter 1

Who Cares About Patent Term? Cross-Industry Differences in Term Sensitivity

1.1 Introduction

There is a debate among scholars over the importance of patent term. Some scholars suggest we move away from our current regime of uniform patent term to one in which term is customized across industry categories (e.g., Budish et al. 2013). They claim patent term might be inadequate for products that take longer to develop, such as certain pharmaceuticals (Roin 2014). Others claim the U.S. patent system is bloated and patent term should be curtailed, or even eliminated altogether, particularly for industries such as software, where patents are thought to be less important (Boldrin and Levine 2008).

But this ongoing disagreement on the importance of patent term glosses over two crucial baseline questions: how do industries differ in how they value term? And why does this variation exist?

This paper addresses both of these questions. First, we exploit a change in law to derive the first cross-industry measure of applicant sensitivity to changes in patent term. Second,

* This chapter has been co-authored with Princeton University Ph.D. alumnus and Associate Professor of Law, Georgetown University, Neel Sukhatme (neel.sukhatme@law.georgetown.edu).

we present a formal model that explains why different industries value term differently: industries in which firms receive more profits toward the end of patent term are more sensitive to changes in patent term. To test this explanation, we then derive an independent measure of late-term patent profits based on cross-industry patent renewal rates. As our model predicts, this measure is highly correlated with our term sensitivity measure.

Our approach throughout is based on a simple idea: an applicant who cares about patent term will seize an opportunity to increase term if it is provided to him. Such an opportunity was created when the United States adopted The Agreement on Trade Related Aspects of Intellectual Property Rights (TRIPS) in 1995. Prior to this law, a patentee received a fixed patent term. After the law, patent term was reduced by the length of time the inventor took to get his patent application approved. We predict applicants who especially cared about term would devote resources to speed up the patenting process after the new law took effect. More importantly, applicants in industries in which patent term is especially important would be more likely to speed up than applicants in industries in which term is less important.

We observe that the amount of applicant speed up varied greatly across industries. In particular, we find support for the conventional wisdom that pharmaceutical applicants are especially sensitive to patent term (Grabowski 2002), as drug companies greatly sped up the patenting process (also known as patent prosecution). We also find that software and computer applicants were more likely to speed up patent prosecution than other groups, which suggests they care more about patent term than previously suggested by scholars (Wagner 2006). On the other hand, mechanical applicants did not speed up prosecution, suggesting patent term is less important to them.

To conduct this analysis, we construct a new dataset based on the transaction histories of 331,162 issued patents filed between 1994 and 1996. Because patent prosecution is a back-and-forth process between applicants and the U.S. Patent and Trademark Office (PTO), we

parse the data to determine what portion of prosecution time is attributable to the patent applicant. Constructing this dataset also allows us to test, for the first time, a prediction that Mark Lemley made 20 years ago: that TRIPS would cause applicants to speed up prosecution (Lemley 1994). Our results confirm his prediction, as we find the average applicant sped up prosecution by about 5.2%.

In addition, to address potential selection issues that might have arisen because applicants changed their behavior in anticipation of TRIPS, we test a number of alternate specifications. For example, we present specifications where we exclude patent applications filed within a certain window around the date the law went into effect. We also test other measures of applicant delay, including the amount of time that applicants took to respond to office actions from the PTO, and the number of extensions that applicants requested during patent prosecution. All our specifications generally produce similar results and identify the same industries as more sensitive to patent term change.

To explain these industry differences in term sensitivity, we present a formal model in which an applicant chooses how long he takes to prosecute his application, taking into account both patent prosecution costs and the timing of expected profits from his patent. We model behavior under two regimes: one in which the applicant obtains a fixed patent term, and one in which patent term decreases with prosecution time. Under the first regime, an applicant will select a prosecution time such that the window in which his patent is enforceable lines up with the window in which he can obtain the most profits, net of prosecution costs. For example, a pharmaceutical company might slow down patent prosecution to align PTO approval with FDA approval of the underlying drug, since the company needs FDA approval before it can sell the drug. Under the second regime, applicants will seek only to minimize prosecution costs, which we model as quadratic.

The model predicts that if a patent office switches from the first regime to the second regime (as the United States did when it adopted TRIPS), applicants no longer have an

incentive to delay patent prosecution. This effect is most pronounced for industries in which higher profits are obtained toward the end of patent term. Put differently, applicants will tend to speed up prosecution in industries in which patents are still profitable at the end of their life. So a testable prediction of our model is that industries in which term sensitivity is high should correlate with industries with higher late-term profits.

We verify the model's prediction using an independent measure of late-term patent profits. In particular, we look across industries and determine the rate at which patentees pay maintenance fees, which a patentee must do to prevent an existing patent from expiring. Maintenance fees provide a (somewhat noisy) measure of the profits generated by patents toward the end of term, so our model predicts that industries in which applicants are more likely to renew patents should be the ones who responded more to the change in law. We confirm this prediction using a difference-in-difference test, finding a high degree of correlation between our term sensitivity and maintenance fee measures. Our results validate our formal model, suggesting that patentees in industries most likely to receive higher profits toward the end of patent term were the ones most sensitive to the change in patent term rules.

1.2 Patent Term and Innovation

In this section, we provide background on patent rights and how TRIPS changed patent term rules in the United States. In the next section, we exploit this change in term rules as part of our empirical strategy to identify the importance of patent term across different industries.

Patents incentivize invention by enabling patentees to monetize and otherwise capture the value of their creations. During the term of a patent, no one can make, use or sell a product containing the invention or perform a method covered by the patent without the

patentee's approval. Without patents, another individual could copy an invention immediately after it became public. Knowing this, the inventor might refrain from producing the invention in the first place, and society would be worse off. So patents can improve social welfare by encouraging the production of new inventions.

Patents also generate a number of costs. These costs certainly include the deadweight loss caused by heightened monopoly pricing. But patents can generate other, less apparent, social costs as well. Too many patents in a technological area increases the risk of inadvertent infringement and can lead to patent thickets, where multiple parties have overlapping sets of patent rights (Shapiro 2001). This can cause holdup problems for new innovators, since those seeking to commercialize a new technology must obtain licenses from multiple patentees (Farrell et al. 2007). Additionally, the uncertainty and high costs associated with patent litigation, even if the underlying patent is weak or actually invalid, can stifle investment and innovation. In particular, suits by patent "trolls" — entities that do not develop technology and acquire patents solely to sue target companies and extract monopoly rents — seem to be a wasteful byproduct of the patent system (Bessen et al. 2012).

One way in which the patent system balances the costs and benefits of patents is by limiting their term. Indeed, Western society has long used patent term to strike this balance. For example, inventors in the ancient Greek city of Sybaris were allowed to collect all profits of any new refinement in luxury, but this right was limited to one year (Anthon 1841). More recently, the U.S. Constitution recognized this balance between incentivizing invention and limiting patent duration, as Article 1, Section 8 authorized Congress, "To promote the Progress of Science and useful Arts, by securing for limited Times to ... Inventors the exclusive Right to their ... Discoveries."

Despite the long tradition of term limits on patent rights, it is unclear what constitutes an optimal patent term. From a social welfare perspective, a patent should be valid only as long as necessary to incentivize invention and no longer, as a longer term only serves to

provide excess monopoly rents to the patentee and incurs the costs described above. And if an invention would be produced even in the absence of patent protection, then it should not receive patent protection at all (Devlin and Sukhatme 2009).

So an important unanswered question is: what set of patent terms optimally balances the benefits and costs associated with patents? And is the current system optimal? There are reasons to doubt this is the case, in the United States or elsewhere.

First, with a few exceptions, all inventions receive the same patent term, regardless of the industry in which they are situated.[1] Whatever one's position might be on what constitutes optimal patent length or whether patents are even necessary, it seems unlikely that very different inventions would optimally receive the same protection for the same amount of time. In some industries, a patent might outlive the usefulness of its underlying invention. For example, in a fast-moving technological field, an invention that is patented might be obsolete well before the end of its term. Such an invention will not yield any profits from sales in the last years of its patented life.

In other technological fields, however, every day of patent term might mean millions of dollars to a company. The paradigmatic example of this is pharmaceuticals — due to the long regulatory approval process before the Food and Drug Administration, a drug may not even be salable for years into its patent term. Indeed, much of the profit for the drug might be obtained in the very last years of its patent term.

So it seems plausible that in an ideal world, the length of patent term would be tailored across different industries. Industries with very high fixed costs, slow obsolescence, or long

[1] Congress has allowed some subcategories of inventions to have slightly longer patent terms, to account for time lost in regulatory approval before the PTO or another federal agency. For example, the Drug Price Competition and Patent Term Restoration Act of 1984 ("Hatch-Waxman Act") allowed certain "pioneer drugs" to receive patent term extensions equal to one-half of the time they spend in the investigational new drug period before the Food and Drug Administration. The American Inventors Protection Act of 1999 (35 U.S.C. § 154(b)) enabled patentees to recover some patent term for applications that take too long to process by the PTO. And in response to lobbying from certain patent owners, Congress occasionally passes bills that extend terms on just these patents.

regulatory delays would feature longer patent terms than industries with lower fixed costs, fast obsolescence, and no regulatory delays.

Second, laws governing patent term are very rarely changed, suggesting that term is not updated in response to changes in technology, society, or the legal system. For example, in the United States, Congress established the patent system in the Patent Act of 1790, which established a term of 14 years from the date a patent issued. Since then, the baseline term for a patent has changed only three times: in 1836 (increased to 21 years from patent issuance), 1861 (decreased to 17 years from patent issuance), and 1995 (changed to 20 years from patent application date). Clearly, many technological, societal, and legal changes have occurred since the founding of the United States, yet the baseline patent term has changed only three times.

Accordingly, it seems likely that the length of patent term under existing law may not be optimal, particularly when comparing terms across industries. To find out whether this is actually true, we need to know the sensitivity of inventors across different industries to changes in patent term.

Unfortunately, few empirical papers have focused on the impact of patent term on innovation, let alone the sensitivity of applicants to changes in patent term.[2] The principal exception is Abrams 2009, which was the first paper to test how changes in patent term caused by TRIPS affected innovation in different industries. Abrams noted that average total pendency (time between patent application filing and patent issuance) was less than three years, with certain patent classes having shorter average pendencies than others. Accordingly, he suggested that these patent classes would benefit disproportionately from the change in law, since they would receive longer patent term extensions on average due to the change in patent term rules caused by TRIPS. He showed that innovation, as measured by

[2] An early effort to measure to the importance of patent term across industries was Levin et al. 1987, which used data from surveys of 650 high-level R&D executives. Their results suggested patents are particularly important in chemical and drug industries, but less so elsewhere.

unweighted and citation-weighted patent application counts, in fields with lower pendency increased disproportionately after TRIPS relative to application counts in fields with longer prosecution times.

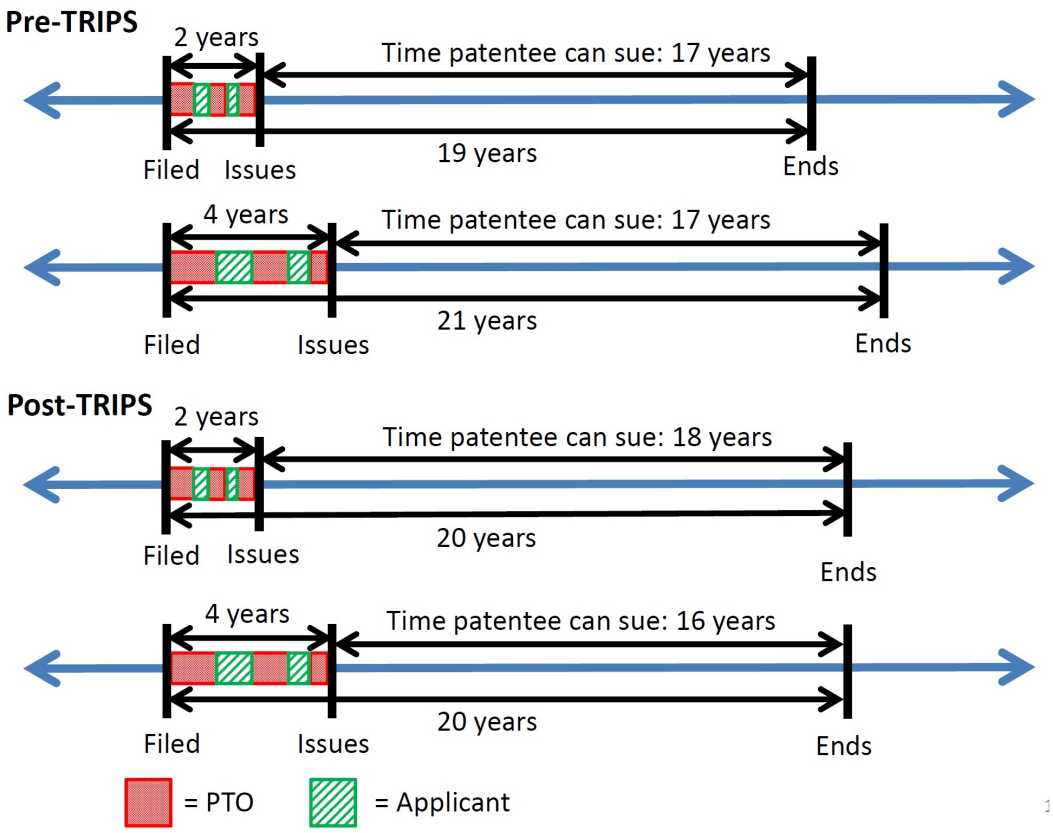
Abrams 2009 set the stage for the present paper. In particular, our paper tests a prediction made by Lemley 1994, which is that after TRIPS, patent applicants would speed up patent prosecution. Prior to TRIPS, patent term was a fixed 17-year term measured from the date on which the PTO granted the patent application. After TRIPS, patent terms were 20 years, measured from the date the inventor applied for the patent. Figure 1.1 on the following page illustrates how patent term changed for patents with two and four-year prosecution times, respectively. Regardless whether patent prosecution took two years or four years, patent term was 17 years before TRIPS went into effect. However, after TRIPS, every day of patent prosecution decreased patent term by one day. So if patent prosecution took two years, then a patentee would have 18 years of effective patent life, whereas if prosecution took four years, then a patentee would have only 16 years of effective patent life.

As Figure 1.1 indicates, patent prosecution is a back-and-forth process between the applicant and the PTO — at some points in time, it is the applicant's turn to act and at others, it is the PTO's. Hence, the time that a patent is pending can be apportioned between the applicant and the PTO.

Lemley's prediction regarding changes in patent prosecution speed has remained untested until now, since there was no way to determine what portion of prosecution time was attributable to an applicant. We are the first paper to test his hypothesis because we are the first to apportion, between the applicant and the PTO, the time from patent application filing to patent issuance.

As discussed in more detail, our results strongly validate Lemley's prediction and suggest that applicants did indeed respond to the change in law by speeding up prosecution. That this change in applicant behavior varied across industries is what allows us to measure cross-

Figure 1.1: Effect of TRIPS on Patent Term



industry patent term sensitivities. In particular, industries in which patent applicants sped up prosecution more can be viewed as more sensitive to patent term; industries in which applicants sped up prosecution less can be viewed as less sensitive.

1.3 Measuring Patent Applicant Delay

We now derive our cross-industry measure of patent term sensitivity. We first provide background on patent prosecution and explain how we constructed our novel dataset. Next, we execute our empirical strategy — namely, we identify how much applicants in different indus-

tries sped up prosecution in response to passage of the TRIPS agreement. We cross-check our results using different measures of applicant delay. We also use a variety of different specifications to address the possibility of selection — that applicants might have changed their behavior in anticipation of TRIPS. Our results remain robust across specifications.

1.3.1 Dataset and Empirical Strategy

Measuring changes in patent applicant behavior requires us to quantify the time an applicant takes to prosecute a patent application. To do this, we use data compiled by Google on the USPTO's Public PAIR (Patent Access Information Retrieval) system (available at: http://www.google.com/googlebooks/uspto-patents-pair.html). PAIR includes public information on all issued patents including their transaction history, which details every substantive or procedural action taken by the applicant or the PTO.

To illustrate, Figure 1.2 on the next page shows a sample transaction history for U.S. Patent No. 5,515,068, entitled "OneChip Semiconductor Integrated Circuit Device Capable of Outputting Analog Color Signal or Digital Color Signal").

This example shows how we might apportion time between the PTO and the patent applicant. After the patent application is filed, the ball is in the PTO's court, as the applicant typically waits for the PTO to either issue a rejection or approve the patent claims.[3] Typically, the first substantive action taken by the PTO is a non-final rejection.

[3]The examiner might also issue something known as a restriction requirement, which is issued when an applicant claims more than one "independent and distinct invention in a single application," and when the patent examiner would face a "serious burden" if forced to examine each of these inventions. In our data, about 10% of patent applications received a restriction requirement at some point during patent prosecution.

In response to a restriction, a patent applicant must file an election in which the applicant chooses which invention to pursue in the present application. Additionally, the applicant may traverse, or dispute the restriction, and he may also pursue the other inventions, if desired, in separate continuation applications. The applicant is generally given one month to respond from the date the restriction requirement is mailed, although the applicant can pay a fee and request an extension of up to five months, giving him up to a total of six months to respond.

We did not see much change in applicant delay in response to restriction requirements before and after TRIPS. It could be that the types of cases in which restrictions were issued were somewhat different from

Figure 1.2: Sample Transaction History

U.S. Patent 5,515,068 - Comp. & Comm. (Category) --> Computer peripherals (Subcategory) --> Computer graphics processing and selective visual display systems (Patent Class)

No.	Date	Event
1	09/02/1994	Preliminary Amendment
2	11/03/1994	Case Docketed to Examiner in GAU
3	11/16/1994	Preexamination Location Change
4	12/12/1994	Non-Final Rejection
5	**12/14/1994**	**Mail Non-Final Rejection**
6	**06/14/1995**	**Request for Extension of Time - Granted**
7	**06/14/1995**	**Response after Non-Final Action**
8	07/16/1995	Date Forwarded to Examiner
9	08/15/1995	Examiner Interview Summary Record (PTOl - 413)
10	08/23/1995	Supplemental Response
11	08/24/1995	Date Forwarded to Examiner
12	10/30/1995	Examiner's Amendment Communication
13	10/30/1995	Mail Examiner's Amendment
14	10/30/1995	Notice of Allowance Data Verification Completed
15	**10/30/1995**	**Mail Notice of Allowance**
16	01/29/1996	Amendment after Notice of Allowance (Rule 312)
17	**01/30/1996**	**Issue Fee Payment Verified**
18	02/21/1996	Date Forwarded to Examiner
19	02/27/1996	Response to Amendment under Rule 312
20	02/27/1996	Mail Response to 312 Amendment (PTO-271)
21	04/01/1996	Issue Notification Mailed
22	04/08/1997	Post Issue Communication - Certificate of Correction
23	05/06/1997	Req. to Make of Record Noted Concerns in Granted Patent
24	05/07/1996	Recordation of Patent Grant Mailed

Here, we can see such a rejection was mailed by the PTO on December 14, 1994. In our data, approximately 77% of patent applications received a non-final rejection like this at some point during patent prosecution.

To avoid abandonment of the patent application, an applicant is required to respond to such a rejection with an argument and/or an amendment to his patent claims, which define his invention. Applicants are typically given three months from the date a non-final rejection

other cases, or that the potential for continuation applications changed the calculus of applicants' behavior. For simplicity, we exclude patent applications involving continuations here and plan to look at this issue more closely in further work.

is mailed to file such a response, though they can pay a fee and request an extension of up to three months, giving them up to a total of six months to respond. Here, we can see the applicant did not respond by March 14, 1995, the three-month deadline. Instead, the applicant requested and received an extension, and he filed his response on June 14, 1995.

This back-and-forth process with the PTO nicely illustrates two potential ways of measuring applicant delay. First, we can measure the amount of time that an applicant takes to respond to a non-final rejection.[4] Second, we can measure whether an applicant has requested an extension when prosecuting an application, and if so, the number of extensions he has requested.

After TRIPS, every day that an applicant takes to respond to a PTO action results in one less day of patent term. Prior to TRIPS, this was not the case. So we would expect to see the average response time to non-final rejections decrease after TRIPS relative to before, with the effect most pronounced in the categories in which inventors are most sensitive to patent term. Relatedly, all else being equal, we should expect to see fewer extensions requested by patent applicants when responding to non-final rejections, since pursuing an extension increases prosecution time and hence decreases patent term.[5]

Returning to the sample transaction history, we can see the patent application here was eventually granted by the PTO. On October 30, 1995, the PTO mailed a notice of allowance to the applicant. A notice of allowance tells an applicant that the PTO has found the

[4]Frakes and Wasserman 2013 recently used a related variable, the average amount of time that the PTO takes to issue its first office action, as a measure of delay by the PTO. Their data was annual summary data, disaggregated at the patent class and entity size level (e.g., average amount of time that the PTO took to issue a first office action for large-entity applicants filing patent applications related to cryptography in 1997), and was obtained via Freedom of Information Act requests.

[5]For present purposes, we are not considering applicant responses to "final" rejections by the PTO. The reason for this is that applicant responses to such rejections can be quite varied and difficult to categorize. In response to final rejections, applicants might file something called a "continuing prosecution application" — essentially, applicants pay an additional fee to get another crack at prosecuting before the examiner. Alternatively, applicants might file an appeal to the Board of Patent Appeals and Interferences. An applicant might also abandon his patent application. Here, we exclude patents involving continuations or appeals (about 10% of our sample). In an extension to our present work, we plan to examine how applicant responses to final rejections, including their propensity to file continuations and appeals, changed as a result of TRIPS.

applicant's claims suitable for granting. Before the patent can issue, however, the applicant must pay an issue fee.[6] An applicant has up to three months to pay an issue fee or else the application is abandoned. After the applicant pays the issue fee, the PTO records payment in the transaction history as "Issue Fee Payment Verified," which happened here on January 30, 1996. Previous empirical work has indicated the patent will then issue, on average, about 6-8 weeks later (Crouch June 9, 2010). For our sample, we find the time between payment of the issue fee and patent issuance is positively correlated (0.1428) with the time between mailing of the notice of allowance and patent issuance. Appendix Figure 1A.1 on page 58 shows the nearly linear relationship between when the issue fee is paid and the average time when the patent issues.[7]

The amount of time that applicants take to pay an issue fee is a third way of measuring applicant sensitivity to patent term. If a patent applicant wants her application to issue quickly after receiving a notice of allowance, she will pay her issue fee quickly. Post-TRIPS, if the applicant cares less about the additional two or three months she might gain by acting quickly, then she might not hurry in paying the fee. In our sample, about 57.15% of the applicants waited at least 80 days to pay the issue fee prior to TRIPS; 54.61% waited at least 80 days after TRIPS.[8] Our theory suggests applicants who care about the few days

[6]The amount of the issue fee has changed over time and varies depending on whether or not an applicant is a "small entity." On March 16, 2013, a "micro" entity class was also created as part of the implementation of the America Invents Act, which was passed by Congress and signed into law in late 2011.

[7]An applicant who plans to enforce her patent would have an incentive to pay her issue fee quickly, even if she does not immediately plan on filing a lawsuit, because damages for patent infringement are typically limited to infringing activities that occur after the patent has issued. So if an infringing activity is already ongoing, a patentee increases the damages she could receive if her patent issues faster.

[8]It is not totally clear why so many people wait until the end of time period to pay the issue fee. One possibility is that applicants must file any continuation patent application — new applications with different patent claims that use the same disclosure as a currently-pending patent application and obtain the same priority date as that application — before the parent patent application issues. By delaying payment of the issue fee, applicants might be increasing the time for filing such continuation applications. As part of our extension regarding continuations, we plan to examine how the timing of issue fee payments might relate to patent continuations.

or months they could gain from acting faster regarding payment of issue fees will speed up their payment of such fees after TRIPS.[9]

Accordingly, our baseline empirical strategy is clear — look at the average time that patent applicants take in responding to non-final rejections and in paying their issue fees, and see if this average decreased after TRIPS. We can also add up both of these types of delay, which is what we do when creating our preferred measure of total applicant delay. Alternatively, we can look at the number of extensions requested post-TRIPS and see if that number decreased. [10]

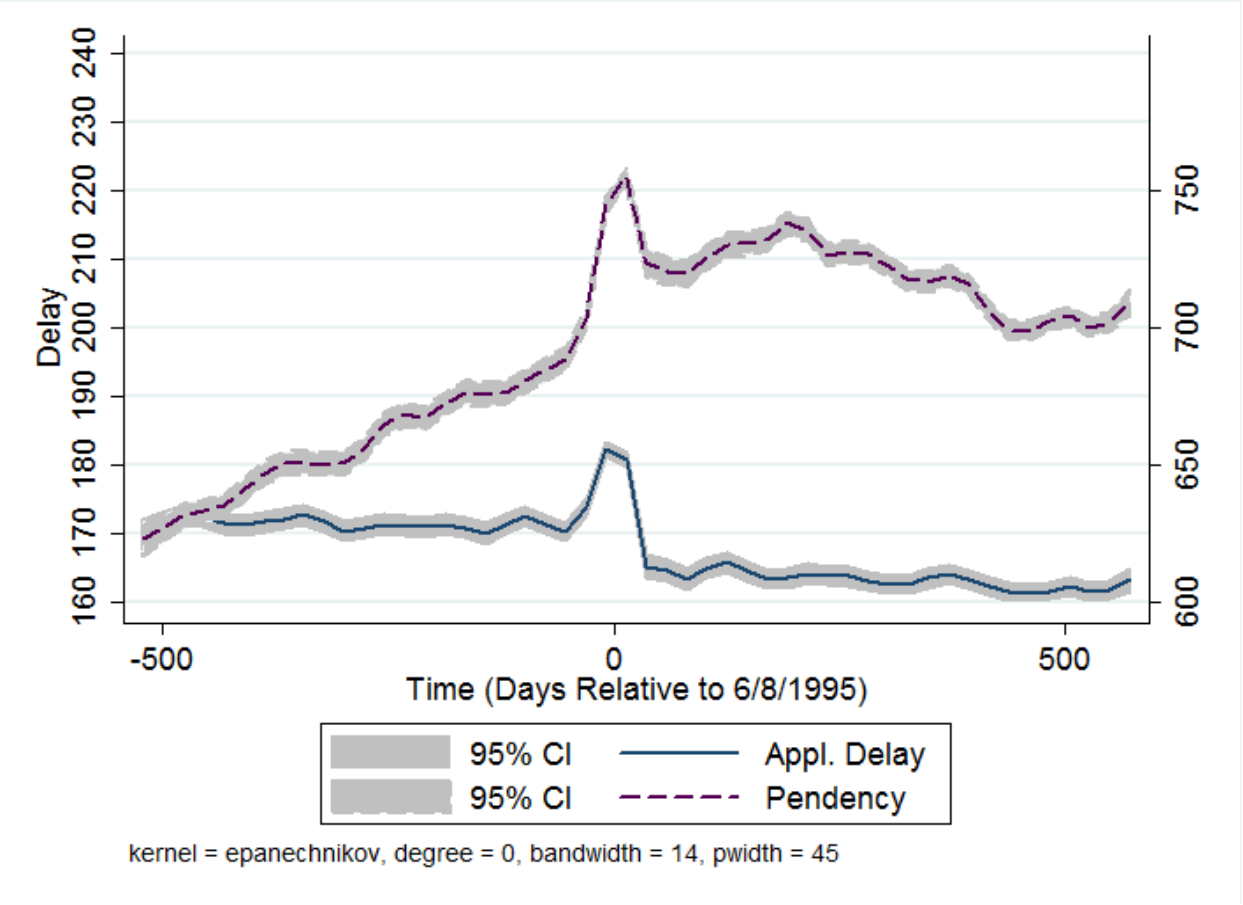
1.3.2 Summary Statistics and Non-Parametric Graphs

Before proceeding to our analysis and results, we show some summary statistics and non-parametric graphs that help describe the data. Figure 1.3 on the next page shows how application pendency — the total time between patent application filing and patent application issuance — and applicant delay — the measure of time attributable to applicants during prosecution that we created using a detailed transaction history for each patent — evolved for patent applications filed between 1994 and 1996.

[9]It may seem somewhat puzzling why applicants have not sped up behavior more in the pharmaceutical context (in the drugs subcategory, 66.55% waited at least 80 days to pay the issue fee pre-TRIPS and 62.85% did the same post-TRIPS). One explanation might relate to the fact that the inventions underlying these patents typically must obtain separate regulatory approval by the Food and Drug Administration. Hence, it may not matter if a pharmaceutical company speeds up patent prosecution, either before or after TRIPS, because it cannot market the product without FDA approval anyway. It is also plausible that continuation applications play a more important role in the pharmaceutical context.

[10]One might note that an applicant may have an incentive to prosecute a patent application quickly even when it does not increase term length. If an applicant cares about receiving a patent sooner rather than later — perhaps because she is trying to obtain a patent as an asset, use it to persuade would-be investors for funding, or simply to advertise or promote the product as "patented" as compared to "patent pending" — then we would expect the applicant to hurry up and not waste time during patent prosecution. This incentive, however, would have existed both before and after the change in law embodied by TRIPS. In other words, if an applicant sped up patent prosecution simply because she wanted to get the patent quickly, not because she cared about patent term, then she would have done so regardless how patent term was calculated. That applicants sped up patent prosecution after TRIPS as compared to before suggests that the real motivator behind the change in behavior was the change in rules regarding patent term.

Figure 1.3: Pendency and Applicant Delay for Applications Filed in 1994-1996



The left y-axis is in days of applicant delay; the right y-axis is in days of pendency. "0" represents June 8, 1995, which is when the new TRIPS regime went into effect. The gray area around each line represents 95% confidence interval bands.

As the graphs show, pendency increased steadily throughout 1994 and into 1995. As June 8, 1995, approaches, we see pendency spike up even more. Meanwhile, our measure of applicant delay also spikes, though to a lesser extent. These trends suggest that increases in pendency during this time were primarily due to changes at the PTO rather than applicant behavior. In particular, the PTO was taking longer to examine applications during this time period.[11]

[11] One possible reason for this trend is that the PTO's workload was growing during this time period because the number of filed applications was increasing steadily in the early 1990s, from 164,306 in 1991 to 189,857 in 1994.

Mindful of the spike in applicant delay, we take seriously the possibility there might be some kind of change in applicant behavior or selection that affects our results. TRIPS was signed into law on December 8, 1994. The law provided that any application filed prior to June 8, 1995, would receive the longer of 17 years from patent issuance and 20 years from patent filing. It is quite possible that some applicants took advantage of this provision by choosing to file patent applications with expected longer pendencies prior to the change in law.

To illustrate, if an applicant had a portfolio of potential applications, he might choose to file the applications that he expects would take longer than three years to prosecute prior to the law change. Such applications might involve, for example, more complicated or pioneering inventions that the applicant expects will require more back-and-forth with the PTO, or they might be inventions in a crowded field with other prior art that makes obtaining a patent a longer and more difficult process. The last months before June 8, 1995, would have also been a good chance for applicants to file "submarine patents" (Lemley 1994). These are patent applications that inventors keep pending in the PTO for a very long time but receive a full 17-year term when they finally issue. This delay might be intentional, so that when the patent finally emerges from prosecution, it enters an already-mature industry in which competitors must now pay to license the technology.

We account for this potential selection effect in our regression analysis by excluding in many specifications those patent applications filed between December 8, 1994, and June 7, 1995; in other specifications, we exclude the entire year 1995. As discussed below, our results remain robust to these changes.

Moreover, rather than viewing selection as something to excise or control for, we could view selection as yet another indicator of patent term sensitivity across industry categories. In other words, the industries in which the most selection occurs are the ones that care the most about patent term. We explore this possibility in more detail in section 1.3.6 below.

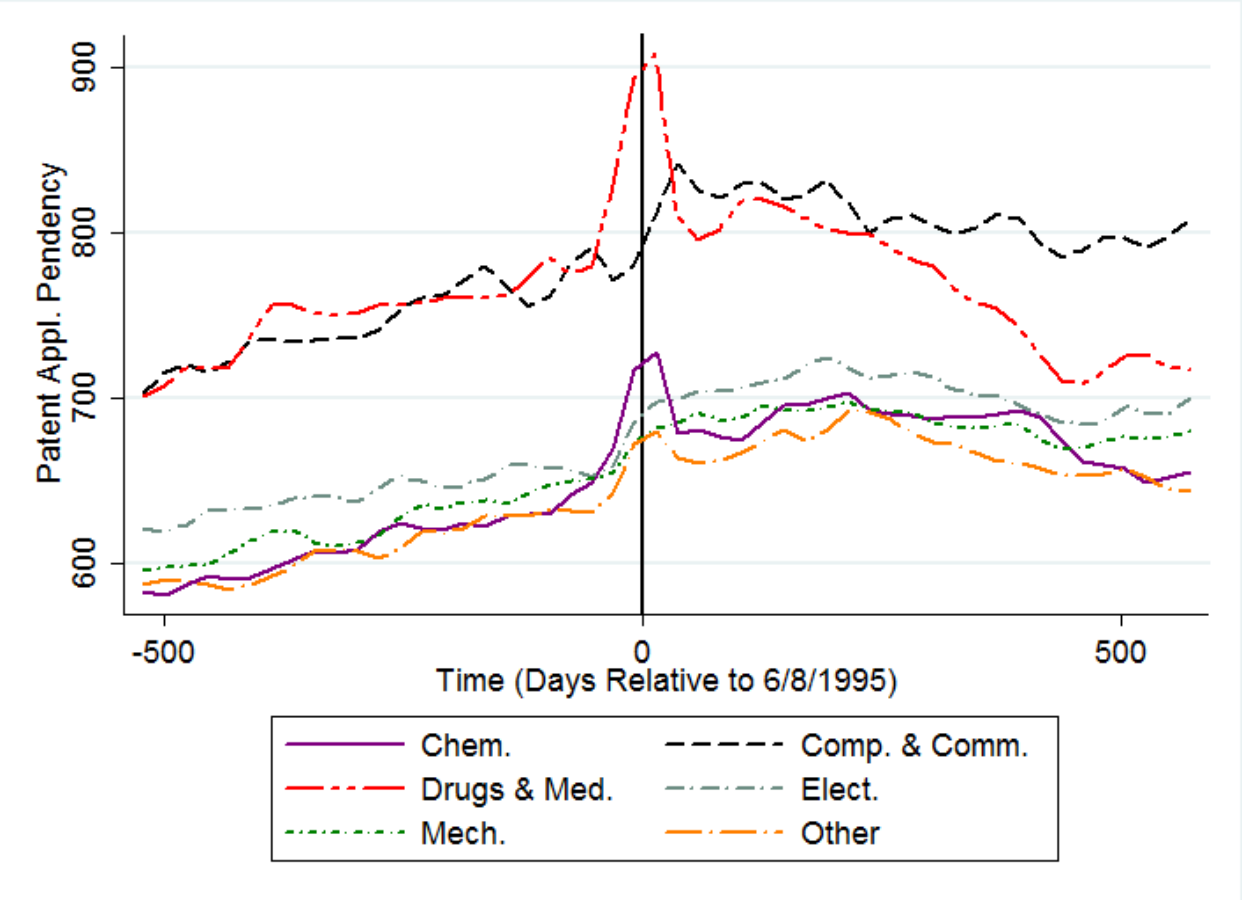
Returning to Figure 5, we see that both pendency and applicant delay drop abruptly after June 8, 1995. Both graphs then become quite stable, with perhaps a slight decrease over time. So the graphs suggest there was a trend break in both pendency and applicant delay around June 8, 1995. Applicants on average took less time to prosecute their applications after the law change as compared to before, in accordance with the theory discussed above. Pendency also switches from increasing steadily to slightly decreasing.

We might be curious whether this trend break only affects certain categories of inventions, or whether its effect is more widespread. Figures 1.4 and 1.5 break down by category changes in applicant delay and pendency, respectively. As we can see, pendency was increasing across all categories prior to the law change. Figure 1.4 shows that pendency appears to decrease slightly across most categories after TRIPS goes into effect. Figure 1.5 shows that applicant delay drops significantly, particularly for the Drugs & Medical, Computer & Communications, and Chemicals categories. The Mechanical and Others categories appear to be the least affected groups, and Electrical falls in the middle. We will discuss these results in more detail in section 1.3.5 below.

We can also test for a trend break due to TRIPS using non-parametric methods. In Figure 1.6 on page 20, we show local polynomials with standard error bands that plot applicant delay relative to time. The polynomial on the left runs from January 1, 1994, to June 7, 1995; the polynomial on the right runs from June 8, 1995, to December 31, 1996. As the figure illustrates, there appears to be a break between the two graphs, where the local polynomial after June 8 is significantly lower than the polynomial before June 8.

The left hand polynomial tips upward prior to the cut date, suggesting some applicants may have chosen to file some especially long pendency applications during this time. But the break appears to be significant even accounting for this transitory increase. Moreover, applicant delay drops precipitously after June 8, suggesting that applicants quickly changed their behavior and sped up patent prosecution under the new regime. Indeed, applicants

Figure 1.4: Pendency by Category



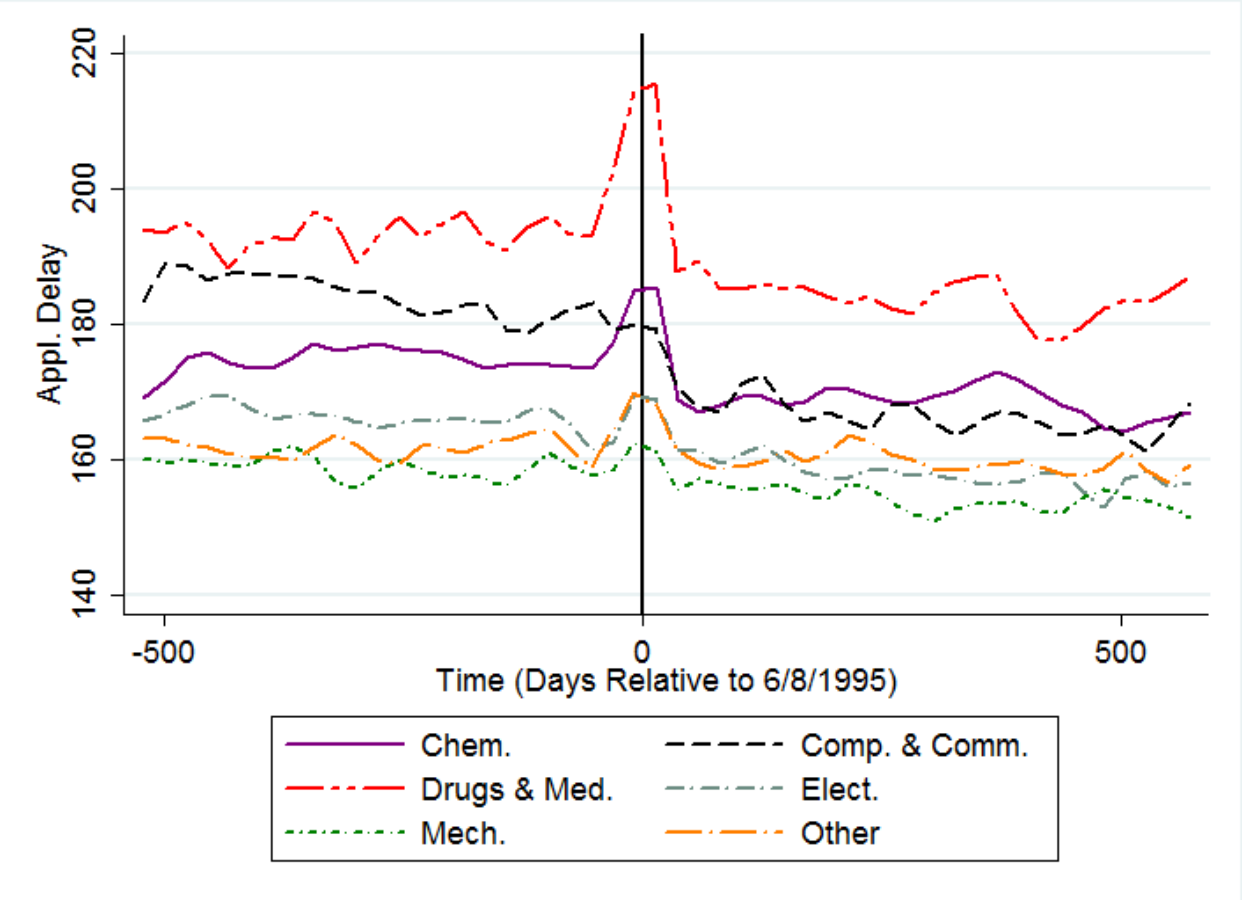
Pendency (total prosecution time) by category using Epanechnikov kernel of degree 0 and bandwidth of 14. "0" represents June 8, 1995, which is when the new TRIPS regime went into effect.

took on average about 9 days less to prosecute their patent applications after TRIPS (drop from about 172 applicant delay days through most of 1994 to about 163 days by the end of 1996).

1.3.3 Results Using Applicant Delay Measure

To validate the suggestive graphical evidence from the previous section, we present a number of ordinary least squares (OLS) regressions that test how applicant behavior changed around the passage of TRIPS. Table 1.1 shows how our preferred measure of applicant behavior,

Figure 1.5: Applicant Delay by Category

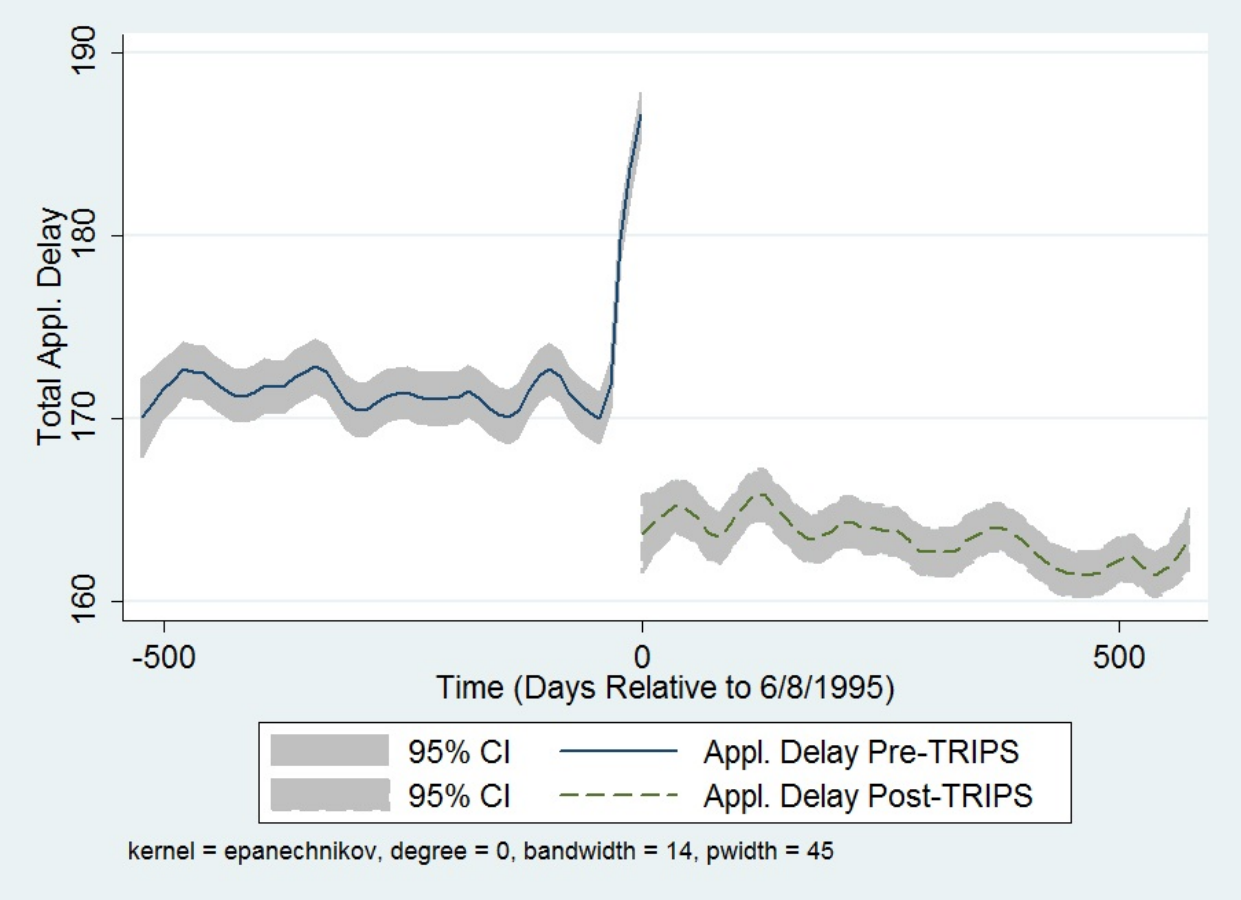


Applicant delay (prosecution time attributable to the applicant) by category using Epanechnikov kernel of degree 0 and bandwidth of 14. "0" represents June 8, 1995, which is when the new TRIPS regime went into effect.

applicant delay, changed significantly around the passage of TRIPS. This is apparent from the coefficient on "Post-TRIPS" — a dummy variable that specifies whether a patent application was filed on or after June 8, 1995, when the new patent term regime went into effect. In our most parsimonious model presented in column (1), the average amount of delay attributable to the applicant fell by about 3.28 days following the implementation of TRIPS. Columns (4)-(6) shows that this result holds with the inclusion of various fixed-effects, as coded by the PTO. These include fixed effects for patent assignee (which is essentially a firm fixed effect),[12] and patent class.

[12] 2.28% of the patents (7,556 of 331,162) have more than one assignee. These patents are excluded in regressions that include assignee fixed effects; results are similar if these assignees are included.

Figure 1.6: Local Polynomial Test for Trend Break in Applicant Delay



Local polynomial test for trend break in applicant delay using Epanechnikov kernel of degree 0 and bandwidth of 14. The left polynomial runs from January 1, 1994, to June 7, 1995, and the right polynomial runs from June 8, 1995, to December 31, 1996. Gray areas show 95% confidence intervals.

Due to possible concerns about selection around the implementation of TRIPS, we also present specifications where we exclude periods during which one might worry most about changed applicant behavior. Our results remain highly significant and of a similar order of magnitude when we exclude either the period between December 8, 1994, and June 7, 1995, (columns (2) and (5)), which is the period between when TRIPS was enacted and when it went into effect, or when we exclude all of the data from 1995 (columns (3) and (6)).

Because an applicant's total delay likely depends on the number of back-and-forth interactions he has with the PTO, we also include controls for number of non-final office actions

and number of restriction requirements for each patent application. The results are robust whether we include these controls or not, or whether we limit to individuals who received at least one non-final office action and/or restriction requirement or not.

Additionally, our results remain similar and highly significant in specifications where we control for the amount of delay attributable to the PTO (remainder pendency in columns (4)-(6)), by subtracting applicant delay from the pendency of the application. In these specifications, we also include controls for whether the applicant is a small entity, whether the patentee was from the United States, and whether the applicant was pro se or instead represented by a patent attorney or agent ("pro se" measures if there was a registered attorney/agent listed for the case). Small entity status, U.S. patentee status, and pro se status seem to have no significant effect on applicant delay. Moreover, including interactions of the various controls described here or including fixed effects for patentee's country of origin do not appear to materially change the results.[13]

One might also be concerned that the above results were driven by a downward trend in our preferred measure of applicant delay. In other words, perhaps applicants were taking less time to prosecute applications as we moved forward in time during the 1990s. If so, the negative coefficients would be a function of this time trend, not because of the change in law.

We can control for this possibility by including time variables in our OLS regressions. These specifications, which we show in Table 1.2 on page 24, help capture any time trend that might have been occurring independent of the change in law. We include a linear variable for time in one model and a linear and quadratic variable for time in a second model. To test for a break in trend, we also include interactions of time and time*time with Post-TRIPS (time*Post-TRIPS and time*time*Post-TRIPS, respectively) in our original

[13]Our results also remain robust in log specifications, where our dependent variable is the log of applicant delay, or when we trim the dataset by excluding applications with very long applicant delays (e.g., where applicant delay exceeds 300 days).

Table 1.1: Effect of TRIPS on Applicant Delay

	(1)	(2)	(3)	(4)	(5)	(6)
Post-TRIPS	-3.28^a	-1.34^a	-1.18^a	-2.25^a	-1.55^a	-1.72^b
	(0.20)	(0.23)	(0.24)	(0.53)	(0.60)	(0.67)
# OAs	106.3^a	105.4^a	105.7^a	104.2^a	103.4^a	103.8^a
	(0.23)	(0.25)	(0.28)	(1.39)	(1.35)	(1.43)
# Restricts.				1.72^a	2.22^a	1.48^a
				(0.42)	(0.48)	(0.55)
Small Entity				1.13	1.10	1.45
				(1.03)	(1.17)	(1.31)
Pro Se				0.63	0.67	0.87
				(0.87)	(1.02)	(1.16)
Remainder Pendency				0.01^a	0.01^a	0.01^a
				(0.001)	(0.001)	(0.001)
U.S. Patentee				-1.27	-1.25	-5.70
				(6.68)	(8.39)	(6.76)
Constant	75.50^a	74.34^a	74.29^a	78.76^a	82.61^a	86.15^a
	(0.21)	(0.25)	(0.26)	(9.87)	(11.61)	(13.51)
Years	1994-96	No inner window	No 1995	1994-96	No inner window	No 1995
Fixed Eff.	No	No	No	Yes	Yes	Yes
Adj. R-sq.	0.58	0.58	0.58	0.68	0.68	0.68
Obs.	331,162	259,977	214,937	268,949	210,282	174,130

Dep. var. is appl. delay. Clustered SEs at patent assignee level in ( ) in all cols. except (1)-(3) (robust SEs). Fixed effs. for patent assignee and patent class. Cols. (2), (5) excl. appls. filed 12/8/94-6/7/95; Cols. (3), (6) excl. appls. filed in 1995. ^{a}P$<$0.01,^{b}P$<$0.05.

models. Now, the statistic of interest is the F-test which tests for equality between the Post-TRIPS variable and its interactions with time, a continuation of trend. We see that the null hypotheses are rejected handily in all three specifications. This outcome also holds for the individual components of applicant delay, helping to allay possible concerns about divergent trends having influenced our results.

Table 1.2: Effect of TRIPS on Applicant Delay with Time Trend Corrections

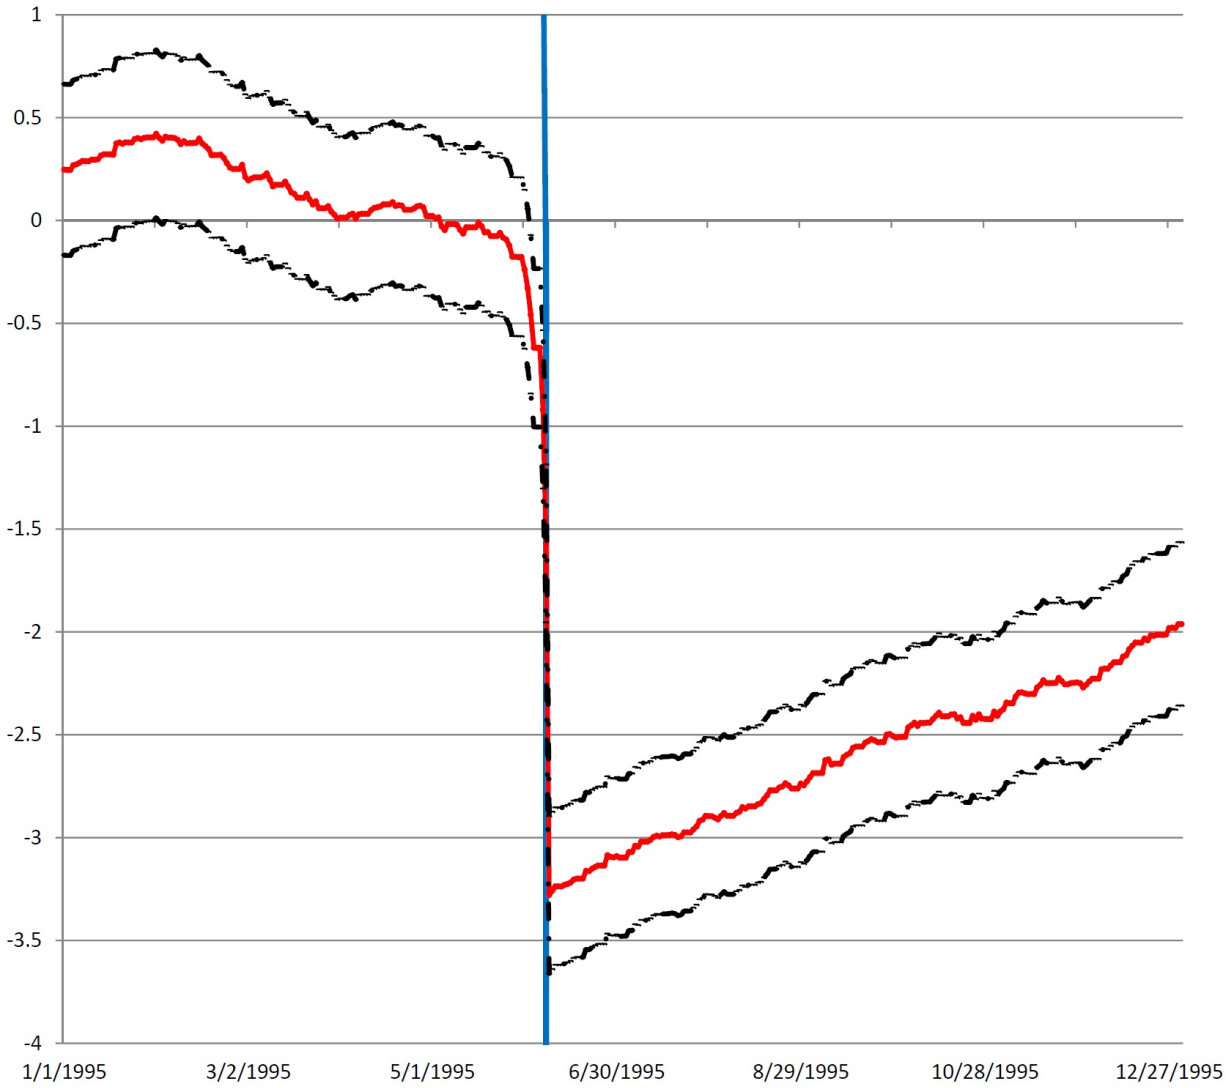
	(1)	(2)	(3)
Post-TRIPS	-1.19	-0.55	2.94
	(0.75)	(1.15)	(4.59)
time*Post-TRIPS	-0.013[a]	-0.005[a]	-0.002
	(0.001)	(0.002)	(0.012)
time*time*Post-TRIPS			-1.95e-05[c]
			(1.11e-05)
time	0.02[a]	0.01[a]	-0.01[b]
	(0.001)	(0.001)	(0.005)
time*time			2.73e-05[a]
			(8.42e-06)
# OAs	106.3[a]	105.0[a]	105.0[a]
	(0.23)	(1.39)	(1.39)
# Restricts.		2.93[a]	2.94[a]
		(0.42)	(0.42)
Small Entity		1.17	1.18
		(1.00)	(1.00)
Pro Se		0.28	0.23
		(0.84)	(0.84)
U.S. Patentee		-0.59	-0.62
		(6.23)	(6.24)
Constant	70.61[a]	79.83[a]	81.42[a]
	(0.33)	(9.59)	(9.59)
Fixed Eff.	No	Yes	Yes
Adj. R-squared	0.58	0.68	0.68
F-Test	282.3	18.21	14.02
Prob > F	**0**	**1.25e-08**	**3.94e-09**
Est. effect on 6/8/95	**-8.09**	**-3.20**	**-3.57**
Observations	331,162	276,397	276,397

Clust. SEs at pat. assignee level in ( ) in cols. (2)-(3); rob. SEs in ( ) in col. (1). Years: 1994-96. F-Test is for Post-TRIPS and its interactions with time. Fixed effs. for pat. assignee and class. [a]P<0.01,[b]P<0.05,[c]P<0.1.

Finally, to help visualize the dramatic impact of TRIPS, we can also use a parametric approach to illustrate the trend break, as shown in Figure 1.7 on the following page. In essence, Figure 1.7 summarizes the results of 365 regressions using the specification in column (1) of Table 1.1. However, instead of using our post-TRIPS dummy variable (which is "1" for patents filed on or after June 8, 1995, and "0" for patents filed before), we slide the dummy variable day-by-day from January 1, 1995, up to December 31, 1995. So, for example, if we choose a dummy around September 28, 1995, the coefficient on the dummy is about -2.5, with a 95% standard error band ranging from about -2.89 to -2.12. This means that, controlling for all covariates, patents filed between September 28, 1995, and the end of our sample period were prosecuted for about 2 1/2 fewer days on average than patents filed from the beginning of our sample period up to September 28, 1995.

If we expect something significant happened on June 8, 1995, we should see a significant drop in our dummy coefficient on that date, and the coefficient would have its lowest value around that time. This is exactly what Figure 1.7 shows — the precipitous decline in the dummy coefficient coincides directly with the blue line marking the start date of the new patent term regime. Applicants spent about 3.28 fewer days prosecuting patents filed on or after June 8, 1995 as compared to before, corresponding with the coefficient on post-TRIPS in column (1) of Table 1.1.

Figure 1.7: Moving Cutpoint Test with Standard Error Bands



Each point represents the coefficient on a dummy variable that is "1" for patents filed after that date and "0" for patents filed before. The specification is according to column (1) of Table 1.1 on page 22, with post-TRIPS replaced by this dummy variable. The vertical blue line shows June 8, 1995, when TRIPS went into effect.

1.3.4 Results Using Disaggregated Measures of Patent Delay

To verify these results are a consequence of faster applicant behavior over the entire process of patent prosecution, Tables 1.3 and 1.4 present the outcome of similar regression analyses on disaggregated measures of applicant delay. Table 1.3 focuses on extension requests

from patent applicants. An extension request (which is nearly always granted) provides an individual extra months to respond to a PTO action, such as a non-final office action or restriction requirement. So cutting down on extensions is a clear way for an applicant to decrease prosecution time.

Looking at the raw data, about 33.9% of pre-TRIPS applicants requested extensions, as compared to 30.4% of post-TRIPS applicants. Columns (1) through (3) present OLS results showing that the average number of extensions fell by about 1.5-2.7 percentage points after TRIPS. Columns (4) through (6) present the marginal effects of logistic regressions on whether an applicant applied for at least one extension. Once again the coefficient on Post-TRIPS is negative and highly significant, suggesting on average that the share of applications with an extension fell by about one percentage point. These results are consistent with a small percentage of patent applicants who changed their behavior due to the change in rules regarding patent term because they highly valued the last years or even months of patent duration.

Table 1.3: Effect of TRIPS on Applicant Requested Extensions

	(1)	(2)	(3)	(4)	(5)	(6)
Post-TRIPS	-0.03[a]	-0.02[a]	-0.02[a]	-0.02[a]	-0.01[a]	-0.01[a]
	(0.002)	(0.004)	(0.005)	(0.002)	(0.002)	(0.003)
# OAs	0.39[a]	0.38[a]	0.37[a]	0.27[a]	0.27[a]	0.26[a]
	(0.002)	(0.015)	(0.015)	(0.001)	(0.002)	(0.005)
# Restricts.	0.21[a]	0.16[a]	0.15[a]	0.10[a]	0.09[a]	0.09[a]
	(0.004)	(0.007)	(0.008)	(0.002)	(0.003)	(0.005)
Small Entity		-0.01	-0.01		0.04[a]	0.03[a]
		(0.010)	(0.013)		(0.003)	(0.004)
Pro Se		0.01	0.01		0.01	0.01
		(0.009)	(0.010)		(0.005)	(0.006)
U.S. Patentee		0.001	0.005		-0.088[a]	-0.099[a]
		(0.074)	(0.081)		(0.002)	(0.005)
Constant	0.032[a]	0.001	-0.065	-	-	-
	(0.002)	(0.089)	(0.082)			
Years	1994-96	1994-96	No inner window	1994-96	No inner window	No inner window
Fixed Eff.	No	Yes	Yes	No	No	Yes
R-squared	0.17	0.30	0.29	0.10	0.11	0.13
Obs.	331,162	268,949	210,282	331,162	215,891	215,881

OLS coeffs. in cols. (1)-(3); dep. var. is # of extension requests. Marg. effects coeffs. for logit in cols. (4)-(6); dep. var. is whether extension request was made. Clustered SEs in ( ) at patent assignee level in cols. (2), (3) and patent class level in col. (6). Robust SEs in ( ) in col. (1). Fixed effs. are for patent assignee and class in cols. (2), (3), and class in col. (6). Adj. R-squared in cols. (1)-(3) and pseudo R-squared in cols. (4)-(6). Cols. (3), (5)-(6) excl. appls. filed 12/8/94-6/7/95. [a]P$<$0.01.

Turning to Table 1.4, we see other measures of applicant delay also decreased on average after TRIPS went into effect. Columns (1) to (3) show that applicants after TRIPS on average took about 1-2 fewer days in responding to non-final office actions, conditional on having received one. Applicants also took about one fewer day paying their issue fee, as

is seen in columns (4) to (6). This result remains highly statistically significant with the inclusion of fixed-effects and across different time periods.

Taken together, these different measures of applicant delay give a consistent picture: applicants appear to have sped up behavior across a number of different facets of patent prosecution in response to TRIPS. These results are robust to the inclusion of fixed effects for patent assignee and patent class of invention, controls for number of non-final office actions and restriction requirements received by a patent applicant as well as small entity, U.S. patentee and pro se status, and various time trends.

Table 1.4: Effect of TRIPS on Non-Final Office Action Response Time and Issue Fee Payment Time

	(1)	(2)	(3)	(4)	(5)	(6)
Post-TRIPS	-2.24^a	-1.35^a	-0.36	-1.55^a	-0.91^b	-0.92^b
	(0.20)	(0.44)	(0.55)	(0.11)	(0.38)	(0.45)
# OAs	107.5^a	105.2^a	103.9^a			
	(0.46)	(1.46)	(1.44)			
Small Entity		-0.29	-0.63		1.50^b	1.79**
		(0.97)	(1.16)		(0.63)	(0.74)
Pro Se		1.33	1.13		-0.53	-0.22
		(0.87)	(0.99)		(0.48)	(0.60)
U.S. Patentee		7.69	8.27		-6.34^b	-6.49^c
		(7.22)	(8.59)		(3.14)	(3.81)
Constant	-0.98^c	-9.17	-7.51	75.17^a	89.34^a	91.46^a
	(0.51)	(6.59)	(7.65)	(0.08)	(7.89)	(9.31)
Years	1994-96	1994-96	No inner window	1994-96	1994-96	No inner window
Fixed Eff.	No	Yes	Yes	No	Yes	Yes
Adj. R-squared	0.47	0.59	0.59	0.001	0.17	0.18
Observations	254,937	207,104	160,936	331,162	268,949	210,282

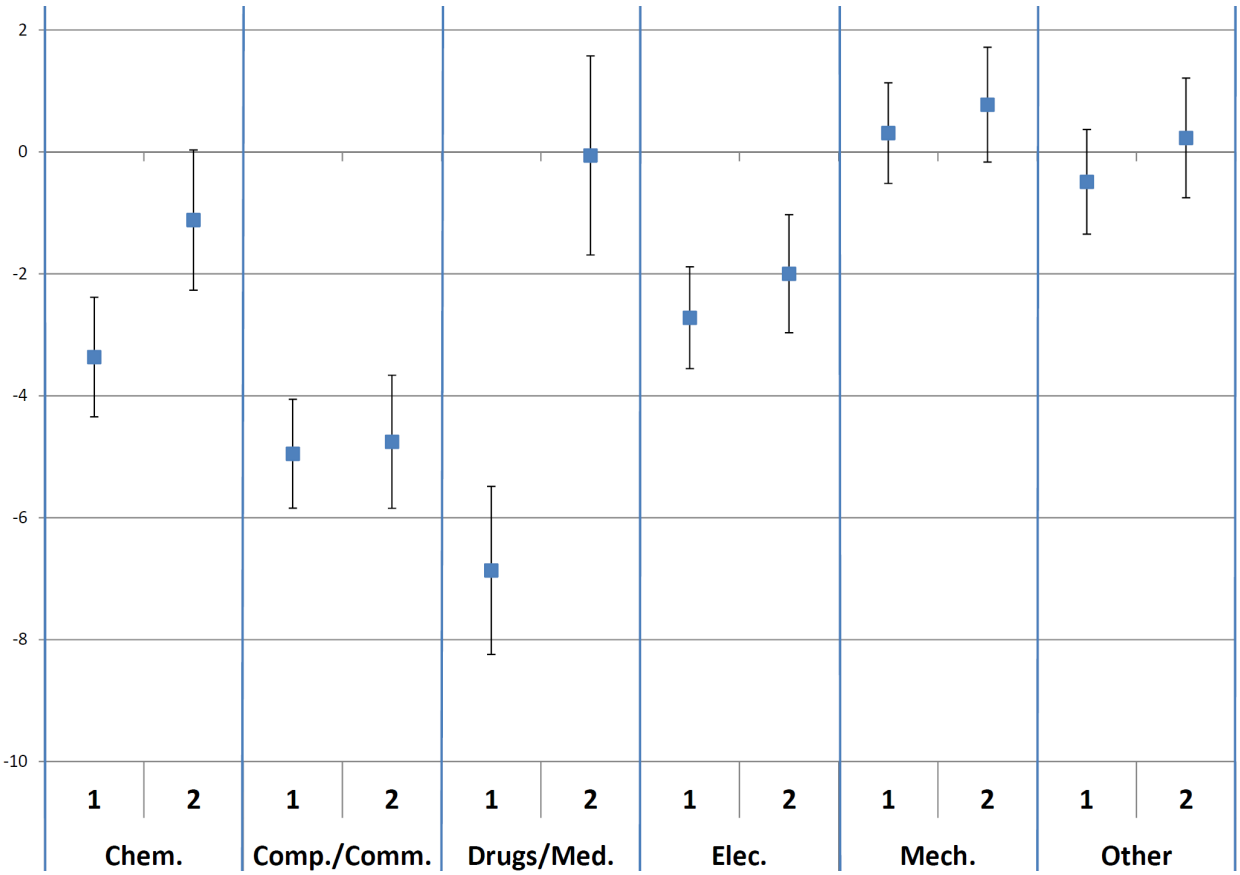
Cols. (1)-(3): dep. var. is total non-final office action resp. time; excl. appls. w/out a non-final rejection. Cols. (4)-(6): dep. var. is time to pay issue fee. Clust. SEs in ( ) at pat. assignee level in all cols. except (1), (4) (robust SEs). Cols. (3), (6): excl. appls. filed 12/8/94-6/7/95. Fixed effs. for patent assignee and class. ^{a}P$<$0.01,^{b}P$<$0.05,^{c}P$<$0.1.

1.3.5 Cross-Industry Differences in Prosecution Delay

We now see how applicants in different industries varied in their response to TRIPS by speeding up their patent prosecution behavior. When assembling the NBER dataset, Hall et al. 2001 consolidated the over 400 classes used by the PTO to categorize inventions into 37 subcategories and six categories. Figure 1.8 on the following page below shows coefficients

for the Post-TRIPS dummy variable with standard error bars for these six categories — Chemicals, Computers & Communication, Drugs & Medical, Electrical, Mechanical, and Others (a residual category). The regression used are our specifications in columns (1) and (2) of Table 1.1 on page 22, where we use applicant delay as our response variable, with a control for number of office actions. The left-hand line for each category is a specification where we use all data from 1994-96; the right-hand line is when we exclude time between when TRIPS was passed and when it went into effect (December, 8, 1994, to June 7, 1995).

Figure 1.8: Cross-Industry Differences in Applicant Delay (Post-TRIPS coeff.)



Graph shows category-level coefficients with standard error bars for the Post-TRIPS dummy variable from specifications in columns (1) and (2) of Table 1.1 on page 22. The left-hand line for each category is a specification where we use all data from 1994-96; the right-hand line is when we exclude time between when TRIPS was passed and when it went into effect (December, 8, 1994, to June 7, 1995).

As Figure 1.8 shows, when using all the data from 1994-96, the largest effect (most significant decrease in applicant delay) appears to be in the Drugs & Medical category and the Computers & Communication category, followed by the Chemicals and Electrical categories. The Mechanical and Others categories appear to have been affected very little by TRIPS, suggesting that applicants in these categories are less sensitive to the change in patent term rules.

When excluding the data from December 8, 1994, to June 7, 1995, we see that the result for Drugs & Medical disappears, but the other results appear similar. This suggests that selection might be particularly important for inventions in the Drugs & Medical category. Patentees in this area might have deliberately filed, prior to the change in law, applications that would likely have a long prosecution time. These applications might take longer because they anticipated a lengthy review from the PTO (perhaps because the inventions were more complicated or difficult to review). Or it is possible that these patentees deliberately delayed prosecution of these applications because they did not want their patent term to begin until later. Such a story would be consistent, for example, with a pharmaceutical patentee deliberately slowing down prosecution of his patent application prior to TRIPS because he still needs to get regulatory approval from the FDA and does not want his patent to issue until such approval has been received. We are exploring these issues in more detail in a follow-up paper.

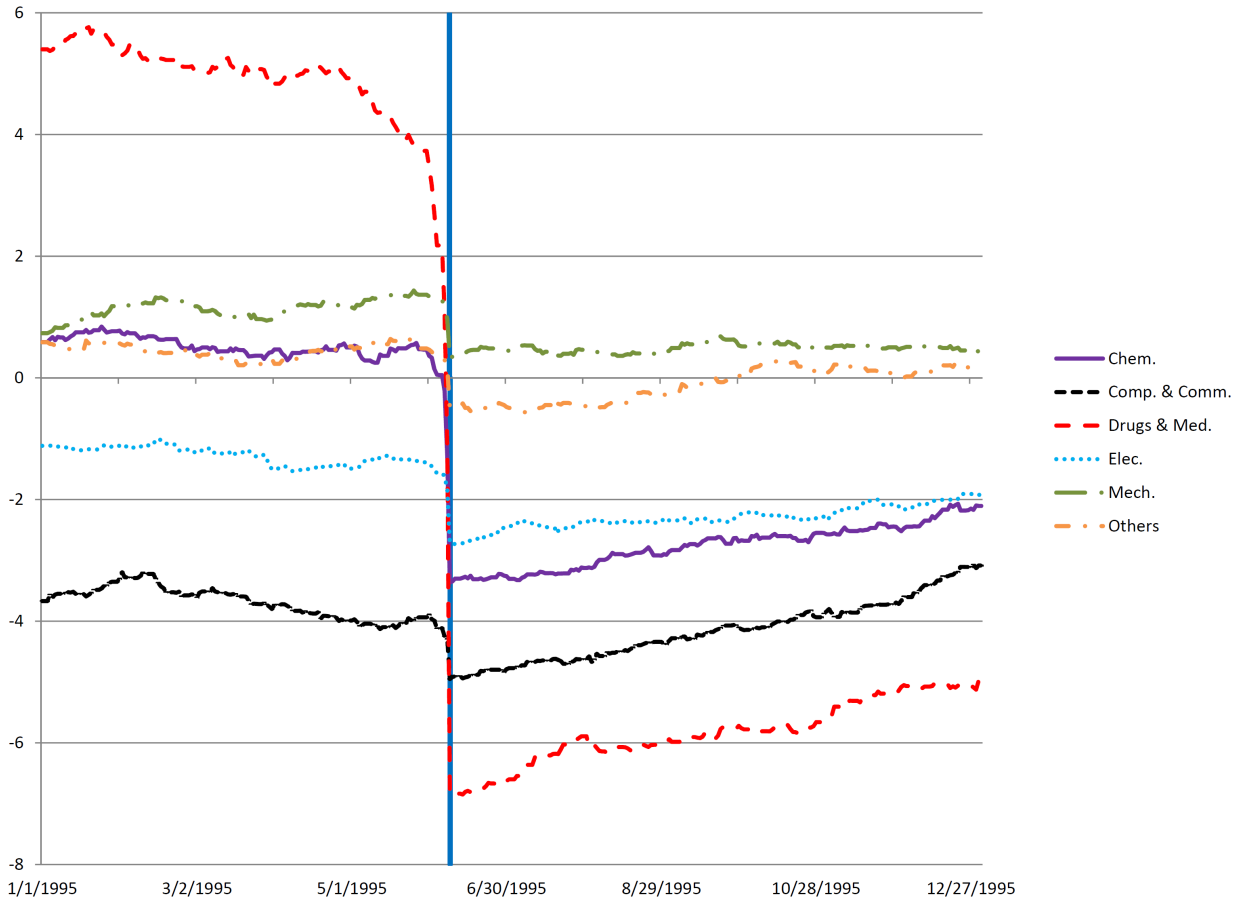
Still, one might have expected the effect for Drugs & Medical to be somewhat stronger, since patents are believed to be important in that industry. A major reason for the results we see here is the way the Drugs & Medical category was constructed. The six categories shown here can be subdivided into 37 subcategories. If one breaks down Drugs & Medical into its component subcategories (see Table 1.6 on page 53, which shows Post-TRIPS coefficients for all subcategories), we see that the subcategories "Drugs" and "Genetics" have very significant, negative coefficients (-2.56 and -18.0, respectively), even when excluding data from December

8, 1994, to June 7, 1995. The coefficient appears to be weakened because "Surgery and Medical Instruments" has been bundled into the Drugs & Medical category, and this is one of the few subcategories with a significant positive Post-TRIPS coefficient (5.17). Accordingly, some of the weaker-than-expected results is attributable to this grouping of pharmaceutical drugs with surgery and medical equipment, which are not thought to rely as much on patent protection.

More generally, heterogeneity among inventions also explains why applicants on the whole sped up prosecution by only a few days after TRIPS — much of the speed-up is concentrated in a few patent classes. Not only is there substantial heterogeneity among subcategories, but each of these subcategories can be further broken down into smaller patent classes, as categorized by the PTO, which show even greater heterogeneity. This can be seen in Appendix Tables 1A.2 and 1A.3, which show Post-TRIPS coefficients significant at the 95% level (under specification 1 in 1.1 on page 22), for patent classes with at least 100 issued patents before and at least 100 issued patents after the change in law. One can see, for example, that the Multicellular Living Organisms, Dynamic Optimal Information Storage or Retrieval, and Organic Compounds classes sped up by 25.75, 14.11 and 12.16 days, respectively, suggesting that applicants in these classes responded much more to the change in law than applicants in other classes.

Figure 1.9 on the next page shows a moving cutpoint test similar to the one discussed above in section 1.3.3, this time with patents separated by invention category. One can see that the effect on the Drugs & Medical coefficient is especially strong. Chemicals also appear to respond strongly, whereas Computers & Communications and Electrical respond less but are negative throughout 1995.

Figure 1.9: Moving Cutpoint Test Across Industry Categories



Patent applications are separated according to the six NBER industry categories. Each point represents the coefficient on a dummy variable that is "1" for patents filed after that date and "0" for patents filed before. The specification is according to column (1) of Table 1.1 on page 22, with post-TRIPS replaced by this dummy variable. The vertical blue line shows June 8, 1995, when TRIPS went into effect.

It is somewhat surprising that TRIPS appears to have had such a strong impact for inventions in the Computers & Communications category. Some of this effect might be due to a time trend — applicant delay in these categories appears to have been trending downward before and after TRIPS. Still, if one runs a time-trend test similar to the one described in section 1.3.3, with inventions divided at the category level, one can see the estimated effect of TRIPS for Computers & Communication remains significant and negative (see Appendix Table 1A.1 on page 59).

1.3.6 Cross-Check Using Patent Count Measure

As noted previously, selection around the date TRIPS went into effect (June 8, 1995), is a potential concern. Applicants who filed their patent applications after that date received a patent term of 20 years from the date of patent application filing. Applicants who filed prior to that date received a term that was the longer of 20 years from the date of patent application filing and 17 years from the date the patent application issued. So if an applicant valued the option of receiving a fixed, 17-year term, then we might expect him to file his patent application before June 8, 1995, if he had some control over when his application would be filed. We dealt with this potential selection by testing specifications that exclude data within certain inner windows, such as applications filed in 1995 or applications filed between December, 8, 1995, (the date TRIPS was passed) and June 8, 1995. Our results have generally remained similar and significant across these specifications.

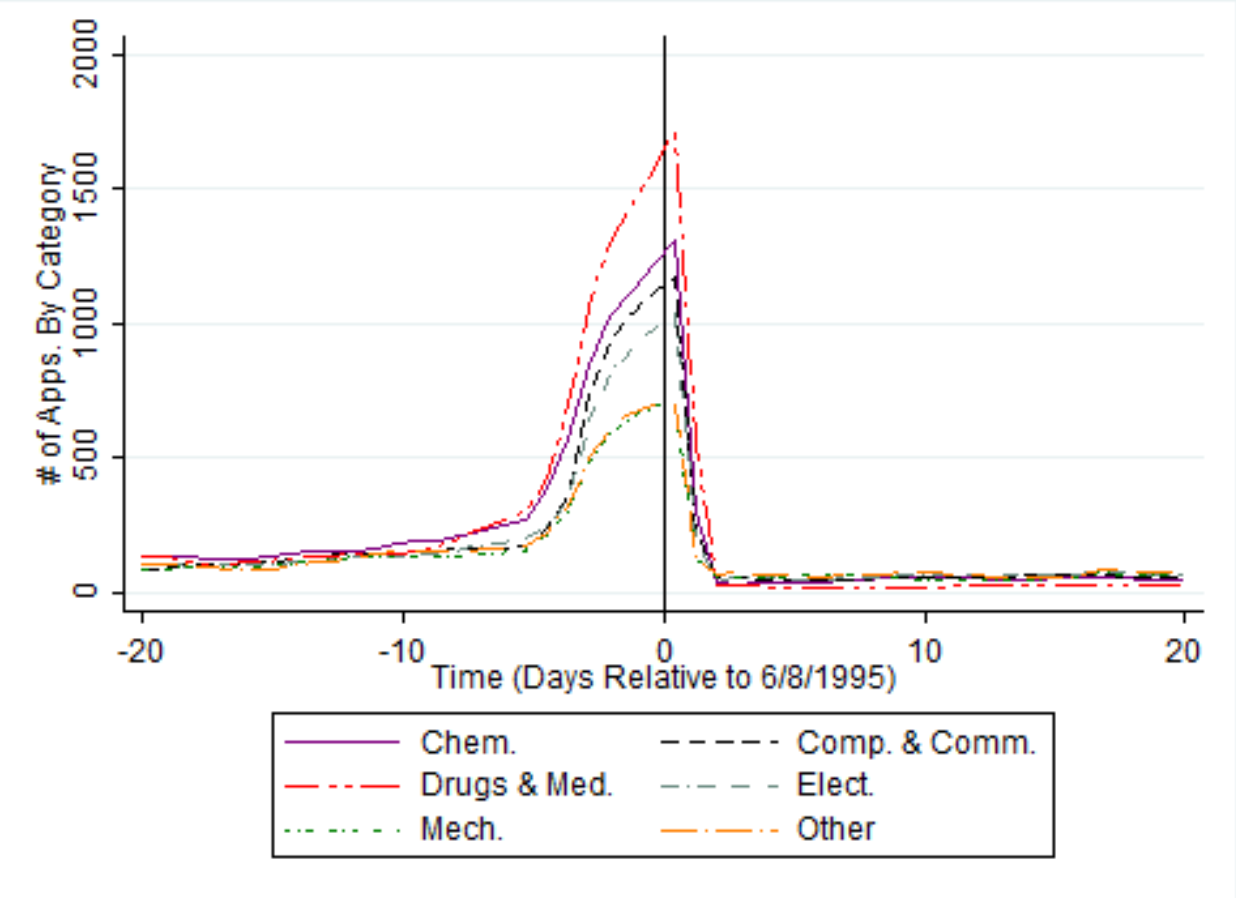
Instead of viewing selection as a potential problem, however, one could instead view it as another measure of applicant sensitivity to patent term.[14] Put differently, it seems plausible that industries in which selection is greater would be the same industries in which applicants care more about patent term. So differences in the level of selection across industries might be indicative of industry-level differences in patent term value. This sort of selection-based measure should correlate with our post-TRIPS measure, which also measures cross-industry differences in patent term sensitivity.

In constructing a selection-based measure, we use patent counts — namely, the number of patent applications filed in a particular industry category within a fixed period of time. One plausible measure of selection is the number of patent applications filed before TRIPS went into effect. As Figure 1.10 on the following page shows, the number of patent applications spiked across all patent categories in the week prior to June 8, 1995, suggesting that

[14] We thank Michael Frakes for helping us with this insight.

many applicants were seeking to file before the new patent term regime took effect. The size of this spike varied greatly across categories — for example, Drugs & Medical applications increased over 17-fold, from about 100 per day to over 1,700, whereas Electrical grew from about 100 per day to over 1,100. This heterogeneity across categories suggests patent applicants in different industries were affected differentially by TRIPS.

Figure 1.10: Number of Applications Filed by Category



Local polynomial regression of number of patent applications filed per day by patent category. Regression uses Epanechnikov kernel of degree 0 and bandwidth of 1, centered on June 8, 1995.

To normalize our patent count measure, we divide the number of patent filings per category by the number of filings from the same time period in the previous year. More formally, our patent count measure is:

$$Count_i = \frac{\#\ of\ patent\ apps.\ in\ category\ i\ (7/1/95-7/7/95)}{\#\ of\ patent\ apps.\ in\ category\ i\ (7/1/94-7/7/94)}$$

To test whether this selection measure is correlated with our primary measure of applicant delay, the post-TRIPS coefficient, we can interact the two measures within difference-in-difference regressions. The results of these regressions are presented in Table 1.5 on the next page, which show that the interaction term between the two measures is negative and significant in all specifications. This means the industries in which applicants sped up prosecution post-TRIPS are the same industries that had the highest percentage increase in patent application filing in the week prior to TRIPS' effective date of June 8, 1995, as compared to filings from a year earlier. The correlation between these measures provides another cross-check that our post-TRIPS coefficient is measuring cross-industry differences in term sensitivity.

Table 1.5: Difference-in-Difference With Count Ratio Measure

	(1)	(2)	(3)	(4)	(5)	(6)
Post-TRIPS	-0.37[a]	-0.18[a]	-0.33[a]	-0.18[a]	-0.18[a]	-0.10[a]
*Count	(0.03)	(0.04)	(0.06)	(0.04)	(0.03)	(0.03)
Post-TRIPS	1.68[a]	0.74	0.03	-3.83[a]	-0.11	-0.38
	(0.39)	(0.47)	(1.02)	(0.98)	(0.46)	(0.55)
Count	1.14[a]	0.95[a]	1.05[a]	0.91[a]	1.37[a]	1.05[a]
	(0.02)	(0.03)	(0.12)	(0.11)	(0.18)	(0.26)
Time			0.011[a]	-0.003		
			(0.002)	(0.002)		
Time			-0.008[a]	0.006[a]		
*Post-TRIPS			(0.002)	(0.002)		
# OAs	105.5[a]	104.7[a]	104.4[a]	103.8[a]	104.2[a]	103.4[a]
	(0.22)	(0.25)	(0.80)	(0.66)	(0.98)	(0.78)
# Restricts.			2.95[a]	3.39[a]	1.73[a]	2.18[a]
			(0.90)	(1.00)	(0.51)	(0.51)
Small Entity			5.56[a]	5.56[a]	1.18	1.36
			(0.57)	(0.64)	(0.98)	(1.22)
Pro Se			-1.34[b]	-0.99	0.62	0.73
			(0.67)	(0.73)	(0.70)	(0.83)
Remainder			0.02[a]	0.01[a]	0.01[a]	0.01[a]
Pendency			(0.002)	(0.001)	(0.001)	(0.001)
U.S. Patentee			-19.94[a]	-20.47[a]	-1.23	-1.32
			(0.86)	(0.79)	(6.10)	(7.36)
Constant	62.80[a]	64.35[a]	64.38[a]	70.22[a]	72.65[a]	77.76[a]
	(0.32)	(0.42)	(2.50)	(2.07)	(4.32)	(6.12)
Years	1994-96	No inner windows	1994-96	No inner windows	1994-96	No inner windows
Fixed Eff.	No	No	No	No	Yes	Yes
Adj. R-sq.	0.58	0.59	0.61	0.61	0.68	0.68
Obs.	331,160	257,959	268,949	208,649	268,949	208,649

Dep. var. is appl. delay. Clust. SEs in ( ) at patent class level in all cols. but (1), (2) (robust SEs). Count = ratio of # patents filed 6/1/95-6/7/95 to # patents filed 6/1/94-6/7/94, w/in patent subcat. Fixed eff. for pat. assignees. Cols. (2), (4)-(5): excl. appls. filed 6/1/94-6/7/94, and 12/8/94-6/7/95. [a]P$<$0.01,[b]P$<$0.05.

1.3.7 Other Changes in TRIPS

Although the change in term rules was the most significant effect of TRIPS on the patent system, we should see whether other TRIPS-initiated changes (Van Horn 1995) might affect our results here. First, TRIPS changed rules relating to establishing an invention date by allowing (for the first time) applicants to establish patent priority based on inventive activity outside of the United States, not just limited to the filing of an international application under the Patent Cooperation Treaty or the filing of a foreign patent application. This seems unlikely to systematically affect the prosecution time of applicants. Moreover, controlling for whether the patentee is from the United States should capture at least some of this effect.

TRIPS also allowed patentees to obtain extensions to their patent term in certain limited situations — when patent issuance is delayed due to an interference proceeding (where two separate inventors dispute who invented first and go through an extensive regulatory process to determine who has priority), a secrecy order (which requires a patent prosecution to remain secret if a government agency believes it to be in the interests of national security), and a successful appellate review. Interferences are very rare — between 1991 and 1994, there were only 718 interferences, much less than 1% of all applications (Calvert and Sofocleous 1995). Secrecy orders appear to be even rarer (at least prior to September 11, 2001), and most of these cases are unlikely to appear in our sample. Regarding appeals, these generally involve very few cases, and even fewer are affirmed (e.g., in fiscal year 1998, there were 3779 appeals received and only 1239 reversed (U.S. Patent & Trademark Office 1998)). At any rate, cases involving appeals have been excluded from our sample, so they do not directly affect our results here.

TRIPS also created a new type of patent application known as a provisional patent application, which is a simplified application (typically without formal claims) that establishes a priority date for the applicant. Provisional applications do not extend patent term; they

simply provide a placeholder for applicants to follow up within one year with a standard, non-provisional application. An applicant who files a provisional and then a follow-up non-provisional within one year gets the priority date of the provisional filing, but receives the standard term of 20 years from the filing of the non-provisional application. Provisional applications that are not followed up with a non-provisional application are abandoned

One concern might be if applicants who systematically differ in delay chose to file provisional instead of non-provisional applications after TRIPS. For example, perhaps applicants who wanted to delay the start of prosecution post-TRIPS filed provisional applications instead of non-provisionals. If we look at only non-provisionals and ignore the provisionals, we introduce a selection bias in our results.

There are reasons to believe this is not an issue here. First, provisional patent applications were not commonly used in the period right after TRIPS. Looking at non-provisional applications filed in 1995-1997, we observe only 2,911 issued patents that claim priority to provisional patent applications filed in 1995 and 6,556 issued patents that claim priority to provisional applications filed in 1996. These are small numbers compared to the 229,471 issued patents filed in 1995-1996 in our sample.

Moreover, our results remain robust when we include in our sample non-provisional applications filed in 1995-1997 that claim priority to provisionals filed in 1995-1996. For example, our Post-TRIPS coefficients for applicant delay remain significant and change from -3.28 to -3.07 for specification 1 of table 1, and -1.72 to -1.66 for specification 6 of table 1. These robustness tests suggest the presence of provisional applications does not materially affect our results.

1.4 Model of Patent Prosecution Time

We now present a formal model, which provides a deeper reason why applicants across industries varied in how much they sped up prosecution post-TRIPS. Our model requires applicants to choose how long they take to prosecute their applications, taking into account both patent prosecution costs and the timing of expected profits from their patent. We model their behavior under two regimes: one in which the applicant obtains a fixed patent term, and one in which patent term decreases with prosecution time. The model predicts that if a patent office switches from the first regime to the second regime (as the United States did when it adopted TRIPS), applicants will tend to speed up prosecution in industries in which more profits are obtained toward the end of patent term. We test this prediction in subsequent sections of this paper.

In the model, a risk-neutral patent applicant i in industry j files an application and chooses the amount of time, $\tau > 0$, she spends prosecuting the application. The applicant's choice of τ depends on two functions: her total cost of prosecuting the patent application, $C_i^j(\tau) > 0$, and the profits she would gain if she has enforceable patent rights at time t, $\pi_i^j(t) \geq 0$. The profits described by $\pi_i^j(t)$ need not be actual profits; rather, they describe the profits an applicant would gain if her patent were issued at or before time t. If the patent has not issued, then those profits remain hypothetical. An example of how such profits might evolve over time is shown in Figure 1.11 on page 43. The figure demonstrates a common situation in which a firm's revenues from patents increase in the beginning, peak after some point, and then decrease in the long-run due to technology obsolescence or some other factor.

Turning to the patent prosecution cost function, $C_i^j(\tau)$, we assume it has an interior minimum; more specifically, the function monotonically decreases from an initial value, $C_i^j(0)$, to a minimum value, $C_i^j(\tau_{min})$, and then strictly increases thereafter for $\tau > \tau_{min}$. This assumption reflects a practical reality for many applicants — total prosecution costs tend to

increase as a patent application is pending, but a patentee who acts very quickly in prosecution also incurs heightened costs. An example of one possible cost function is shown in Figure 1.12 on the next page.

Figure 1.11: Example of Patent Profits Evolution Over Time

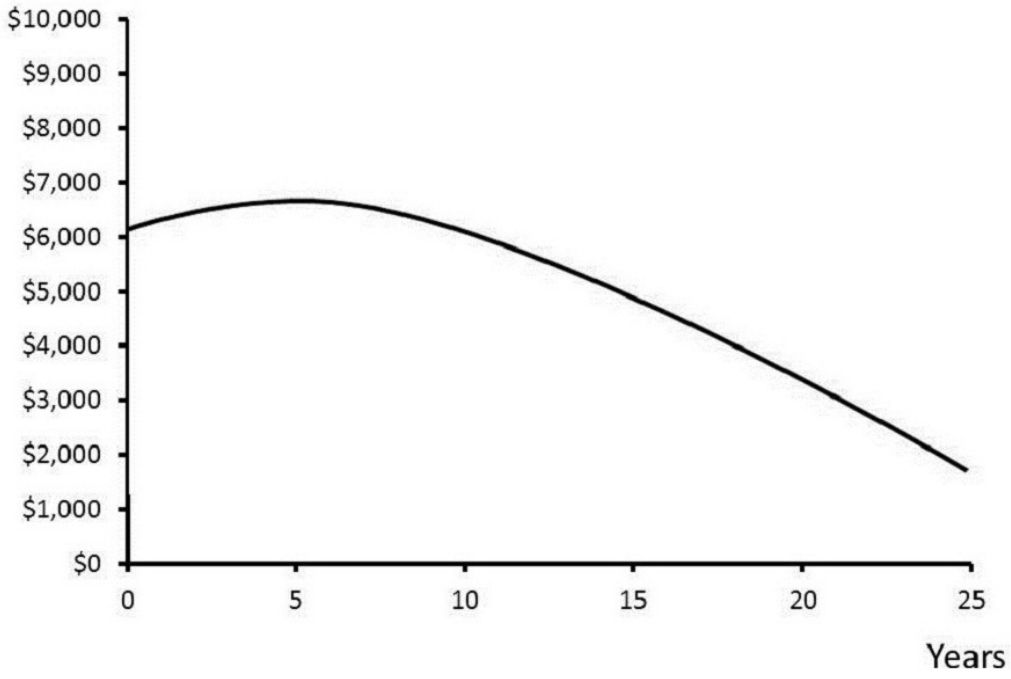
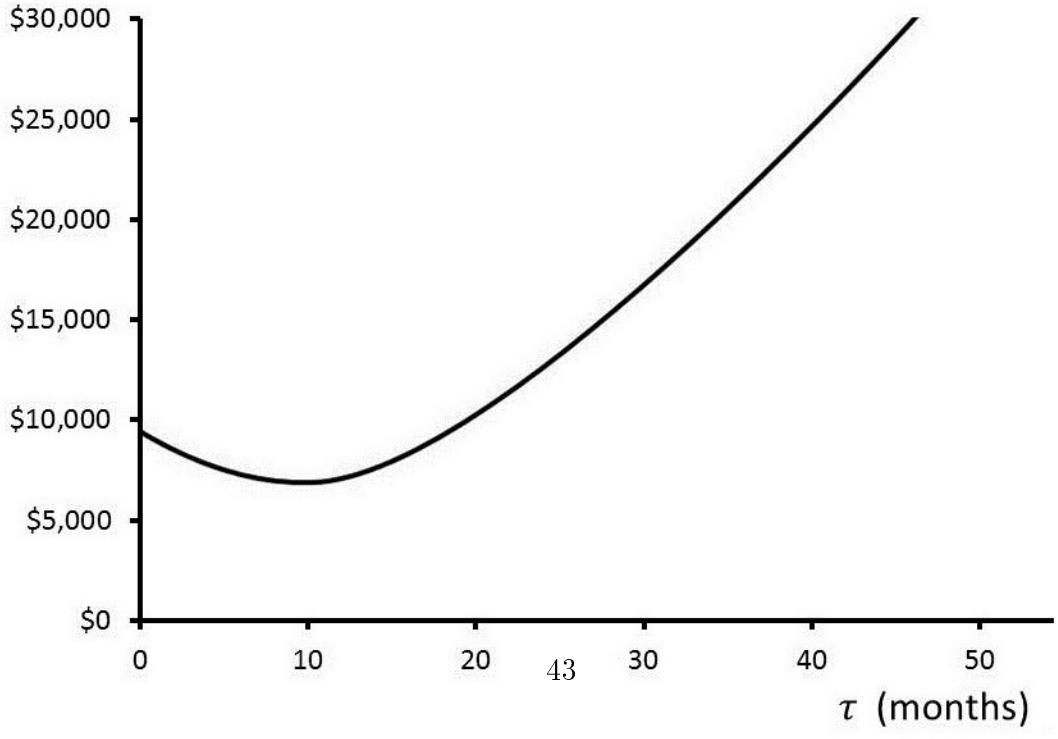


Figure 1.12: Example of Prosecution Cost and Prosecution Time Relationship



Additionally, the firm's optimization decision will depend on $\tau^j_{i,PTO}$, which is the applicant's ex ante expectation of the amount of time the patent office will spend reviewing the application. The applicant's expectation can depend on the particular application at issue as well as the underlying industry for the invention. For simplicity, we will often denote $\pi^j_i(t)$, $C^j_i(\tau)$, and $\tau^j_{i,PTO}$ from here on without superscripts or subscripts, since we are focused on a particular applicant in a particular industry choosing her prosecution time. But one should note these can be flexible functions that vary across applicants and industries.

1.4.1 Model with Fixed Patent Term

Suppose an applicant faces a fixed patent term that begins on the date her patent is granted. The patent term is the same regardless of the time the applicant spends prosecuting the patent application. This was how things were in the United States prior to TRIPS, when patent term was always 17 years as measured from the date of patent issuance. More generally, we denote this patent term duration as T_{old}.

Under this patent term regime, $\tau + \tau_{PTO}$ is the patent grant date; hence, profits can be obtained from that date until patent expiration at $\tau + \tau_{PTO} + T_{old}$. The applicant chooses an optimal τ, denoted τ^*_{old}, based on the following optimization problem:

$$\tau^*_{old} = arg\,max_\tau \int_{\tau+\tau_{PTO}}^{\tau+\tau_{PTO}+T_{old}} \pi(t)dt - C(\tau) \tag{1}$$

We assume $\int_{\tau+\tau_{PTO}}^{\tau+\tau_{PTO}+T_{old}} \pi(t)dt - C(\tau) \geq 0$ for some value of τ, otherwise the applicant would not file a patent application in the first place. We also assume this quantity satisfies all the other normal conditions to ensure that an interior solution exists. If we take the first

order condition and solve we see:

$$\pi(\tau^*_{old} + \tau_{PTO} + T_{old}) = \pi(\tau^*_{old} + \tau_{PTO}) + \frac{dC(\tau)}{d\tau} \mid_{\tau=\tau^*_{old}} \tag{2}$$

This means a patent applicant will choose τ^*_{old} such that the marginal gain of increasing τ^*_{old} — the increased profit from the end of patent term, $\pi(\tau^*_{old} + \tau_{PTO} + T_{old})$ — equals the marginal cost of increasing τ^*_{old} — the lost profit from the beginning of the patent term, $\pi(\tau^*_{old} + \tau_{PTO})$, together with the change in prosecution costs, $\frac{dC(\tau)}{d\tau} \mid_{\tau=\tau^*_{old}}$.

1.4.2 Model with Term Dependent on Prosecution Time

Now suppose patent term depends on the amount of time an applicant spends in patent prosecution. In particular, assume that a patentee receives a fixed term, minus the time that the patentee took to receive the patent application. This became the situation in the United States after TRIPS went into effect, and a 20 year patent term was instituted that began from the date the patent application was filed. Under this patent term regime, as prosecution time τ increases, effective patent term decreases, since the beginning of the fixed patent term is measured from the application filing date.

Under the new patent term regime, the patent grant date remains $\tau + \tau_{PTO}$, but patent profits terminate at T_{PTO}. The applicant chooses an optimal τ, denoted τ^*_{new}, that solves the following problem

$$\tau^*_{new} = arg\,max_\tau \int_{\tau+\tau_{PTO}}^{T_{new}} \pi(t)dt - C(\tau) \tag{3}$$

Solving for the first order condition we see

$$-\pi(\tau^*_{new} + \tau_{PTO}) = \frac{dC(\tau)}{d\tau} \mid_{\tau=\tau^*_{new}} \tag{4}$$

Recall that τ_{min} is the point where $C(\tau_{min})$ is at its minimum. We can use equation 4 on the preceding page to prove the following lemma (proof in Appendix section B.1)

Lemma 1: $\tau^*_{new} \leq \tau_{min}$

1.4.3 Differences in Prosecution Time Under the Two Patent Term Models

Next, we see how an applicant's behavior varies depending on the patent term regime she faces. If we compare the optimal prosecution time of a patent applicant who faces the old patent term regime (where patent term is measured from the date of patent issuance) as compared to her optimal prosecution time under the new patent term regime (where patent term is measured from the date of patent application filing), we can see the applicant never strictly prefers choosing τ^*_{old}, τ^*_{new} such that $\tau^*_{new} > \tau^*_{old}$. At most, and only under certain limited circumstances, will an applicant be indifferent between τ^*_{old} and $\tau^*_{new} > \tau^*_{old}$. If these conditions are not met, then $\tau^*_{new} \leq \tau^*_{old}$ — in other words, the applicant will choose at least as fast a prosecution speed under the new patent term regime vis a vis the old patent term regime.

The preceding statements are stated more precisely in the following two propositions (proofs are in Appendix, sections B.2 and B.3).

Proposition 1: If $\int_{\tau^*_{old}+\tau_{PTO}+T_{old}}^{\tau^*_{new}+\tau_{PTO}+T_{old}} \pi(t)dt \neq 0$, then $\tau^*_{new} \leq min(\tau^*_{old}, \tau^*_{min})$.

Proposition 2: If $\int_{\tau^*_{old}+\tau_{PTO}+T_{old}}^{\tau^*_{new}+\tau_{PTO}+T_{old}} \pi(t)dt = 0$, then:

(1) $\int_{\tau^*_{old}+\tau_{PTO}}^{\tau^*_{new}+\tau_{PTO}} \pi(t)dt = C(\tau^*_{old}) - C(\tau^*_{new})$, and

(2) applicant is indifferent between τ^*_{old}, τ^*_{new} under either term regime.

So we can see that ceteris paribus, an applicant placed in the new patent term regime will spend less time prosecuting his patent application than if the same applicant were placed in the old patent regime, unless there are no additional profits to be made toward the end of patent term by speeding up prosecution under the new regime (i.e., $\int_{\tau_{old}+\tau_{PTO}+T_{old}}^{\tau_{new}+\tau_{PTO}+T_{old}} \pi(t)dt = 0$). An applicant in such a situation — for example, an applicant whose patent will have no value toward the end of its term — will not have the same incentive to speed up patent prosecution, so a switch from the old patent term regime to the new patent term regime may not result in faster prosecution.

If we are willing to specify functional forms, we can also determine the extent to which an applicant will speed up patent prosecution. We assume $C(\tau)$, which represents the applicant's total cost of getting a patent through the patent office, has the following flexible quadratic form:

$$C_i^j(\tau; \alpha_0^i, \alpha_{1,}^i s_j) = \alpha_0^i + \alpha_1^i \tau + s_j \tau^2$$

We assume that $C_i^j(\tau)$ is a function of τ as well as three parameters: an industry-specific fixed cost, s_j, and two firm-specific parameters that influence how costs evolve over time, α_0^i and α_1^i. The industry-specific component, s_j, captures the fact that patenting in certain industries requires more prosecution time than patenting in others, perhaps due to heightened technical complexity or the prevalence of prior art that makes obtaining a patent more difficult. The firm-specific parameters capture firm-specific behavior that might influence the level and evolution of patent prosecution costs. We assume the industry-specific fixed

cost $s_j > 0$ and individual-specific parameters $\alpha_0^i > 0$ and $\alpha_1^i < 0$. Consistent with the more general case discussed above (and Figure 1.12 on page 43), this functional form models the practical reality that while patent prosecution costs tend to increase over time, they can also increase if a patentee tries to prosecute an application too quickly. As before, prosecution costs are minimized at some $\tau_{min} > 0$.

Next, we plug in this functional form into equations 2 and 4 on page 45 and obtain the following:

$$\pi_i^j(\tau_{old}^* + \tau_{PTO} + T_{old}) - \pi_i^j(\tau_{old}^* + \tau_{PTO}) = \alpha_1^i + 2s_j\tau_{old}^* \tag{5}$$

$$-\pi_i^j(\tau_{new}^* + \tau_{PTO}) = \alpha_1^i + 2s_j\tau_{new}^* \tag{6}$$

Together equations 5 and 6 imply:

$$\tau_{new}^* = \tau_{old}^* - \left(\frac{\pi_i^j(\tau_{old}^* + \tau_{PTO} + T_{old}) + \pi_i^j(\tau_{new}^* + \tau_{PTO}) - \pi_i^j(\tau_{old}^* + \tau_{PTO})}{2s_j} \right)$$

Accordingly, the extent to which a patent applicant speeds up prosecution in the new patent term regime relative to the old patent term regime depends both on profits toward the end of patent term and profits at the beginning of the patent term. To the extent profits toward the end of patent term are high (e.g., $\pi_i^j(\tau_{old}^* + \tau_{PTO} + T_{old})$ is large), the applicant is more likely to speed up prosecution. If the difference in additional initial profits from speeding up is also high (e.g., $\pi_i^j(\tau_{new}^* + \tau_{PTO}) - \pi_i^j(\tau_{old}^* + \tau_{PTO})$), then the applicant is also likely to speed up prosecution. Finally, we see that an increase in s_j — which is the strictly positive quadratic term in the cost function that determines how quickly prosecution costs increase as a function of time — will decrease the rate at which applicants speed up prosecution under the new patent term regime relative to the old regime.

It is important to note this model is flexible and allows for the post-TRIPS "speed up" to actually be an elimination of a prosecution delay that existed before TRIPS. More specifically, TRIPS might have eliminated a previously existing incentive for some applicants to delay prosecution of their patent applications. For example, prior to TRIPS, if a potential pharmaceutical patent applicant did not anticipate receiving FDA approval for a drug for nine years, it might delay patent prosecution in order to try to align the start of its patent term with its FDA approval date (since its patent term would essentially be wasted if it also does not have approval of the drug by the Food and Drug Administration).[15]

Regardless whether what happens post-TRIPS is a speed up in applicant prosecution time (because an applicant cares about patent term), or the elimination of a pre-TRIPS prosecution delay (because the applicant previously wanted to delay patent issuance), we can see an applicant will speed up prosecution post-TRIPS only if he expected to receive at least some profits 17 years or more after filing the patent application. Moreover, the amount of prosecution speed-up depends in part on the profits an applicant stands to gain at the end of her patent term. So comparing the prosecution speed of applicants within an industry before and after TRIPS gives us a measure of how much length of patent term matters in that particular industry.

Accordingly, this model provides us with some testable predictions in terms of applicants' response to TRIPS. We should expect that on average, patent applicants will speed up prosecution after TRIPS relative to before. And ceteris paribus, applicants in industries that yield more profits toward the end of patent term (e.g., the right hand tail of the profit function shown in Figure 1.11 on page 43 is thicker) will speed up prosecution more than applicants

[15]Before TRIPS, such an applicant would also have an incentive to delay the filing the patent application, but this strategy would carry more risk — namely, another inventor might seek patent protection on the same invention in the meantime, making it more complicated as a practical matter to prove that the original inventor deserves patent protection as the "first to invent." So prior to TRIPS, slowing down patent prosecution would be a better strategy for pharmaceutical patentees as compared to delaying the patent application filing. This is an issue we are exploring in a follow-up paper: whether pharmaceutical applicants indeed deliberately slowed down patent prosecution prior to TRIPS, so that their patent grant date would be more aligned with FDA approval of the underlying drug.

in industries where inventions lose value quicker and patents are relatively valueless toward the end of term. The reason for this is intuitive: industries where patents generate more revenue toward the end of term have more to lose by decreases in patent term; hence, they will be more likely to speed up patent prosecution in response to TRIPS.

1.5 Maintenance Fees as a Measure of Late-Term Patent Value

In this section, we derive a cross-industry measure of late-term patent value by measuring the proportion of patents within invention categories that do not lapse for failure to pay patent maintenance fees (also known as patent renewal fees). We calculate this measure in order to test our patenting model, which connects two values — the amount of speed-up in prosecution behavior due to TRIPS, and the amount of profits receivable at the end of patent term. Our model suggests that industries in which applicants sped up the most post-TRIPS will be more likely to be the industries in which patents yield the most profits toward the end of term.[16] We confirm our model's prediction through a difference-in-difference test, we shows our measure of applicant delay is highly correlated with our maintenance fee measure.

Since 1981, the PTO has required applicants to pay three maintenance fees, due at $\mathbf{3\frac{1}{2}}$, $\mathbf{7\frac{1}{2}}$, and $\mathbf{11\frac{1}{2}}$ years after patent issuance. If a patentee does not pay these fees, then his patent will lapse, which means its term will end immediately rather than at its scheduled time. The amount of maintenance fees is uniform across industry categories; in 1995, it was \$650, \$1,310, and \$1,980, for $\mathbf{3\frac{1}{2}}$, $\mathbf{7\frac{1}{2}}$ and $\mathbf{11\frac{1}{2}}$-year maintenance fees, respectively.[17] 35 U.S.C.

[16]This correlation may not be perfect, as the model also suggests that profits obtainable at the beginning of patent term also affect the amount of applicant speed-up. But we know that applicants will speed up prosecution only if there are at least some expected late-term profits, which is a testable prediction of our model.

[17]In a follow-up paper (Sukhatme 2014), one of the authors explains how the PTO could incentivize innovation and improve social welfare by charging different maintenance fees to applicants in different in-

§ 41 (1995). If a patent assignee qualifies as a "small entity" — that is, if it is an individual, or a small business or non-profit organization that meet certain criteria — then these fees are halved. About 25% of patents issued between 1981 and 2000 qualified for "small entity" status in our dataset, which we obtained online from Google.

Scholars have long used data on maintenance fee payments as a measure of patent valuation. (Lanjouw et al. 1998, Schankerman 1998, Allison et al. 2003, Moore 2005, Bessen 2008). The idea is that if a patentee chooses to pay a maintenance fee, we can infer the patentee believes his expected benefit from keeping the patent alive exceeds the present cost of the maintenance fee. More specifically, maintenance fees tell us something about discounted expected profits from patents toward the end of their term. For example, a patentee should renew a patent at $\mathbf{11\frac{1}{2}}$ years only if the discounted expected value of future profits stemming from the patent exceeds the maintenance fee payment at that time.

Accordingly, we can use maintenance fee payments as a rough measure of the relative value of patents across industries, particularly late-term patent value. For example, if pharmaceutical patents are more likely to be maintained on average than mechanical patents, then we might conclude that pharmaceutical patents are on average more valuable than mechanical patents. Of course, this conclusion only holds true for patents near the top end of the patent value distribution — that is, those patents whose value exceeds the maintenance fee threshold amount. It might not hold true for less valuable patents within a patent category. In short, maintenance fees cannot tell us the precise distribution of patent values within industries (Abrams et al. April 8, 2013). Nonetheless, a higher rate of renewal within an industry category, particularly at the $\mathbf{11\frac{1}{2}}$ year date, suggests that on average, patents in that industry are worth more toward the end of their term than in industries with lower renewal rates. So if our model is correct, we might expect to see some correlation

dustries, using the cross-industry sensitivities identified in this paper. Such pricing would take advantage of the PTO's inherent monopoly in the pricing of user fees as well as its new power to tailor such fees, which was granted as part of recently-enacted America Invents Act. .

between industries that have higher renewal rates and industries in which applicants sped up prosecution the most in response to TRIPS.

The data in Table 1.6 on the following page shows the percentage of patents filed between 1994 and 1996 that were maintained through $\mathbf{11\frac{1}{2}}$ years within the 37 subcategories described above. It also shows the post-TRIPS coefficients using our specifications in columns (1) and (2) of Table 1.1 on page 9.

Even though there are only 37 subcategories, we see a strong and significant negative correlation between the maintenance fee measure and our measures of applicant delay (e.g., -0.482 and -0.516 correlation, significant at the 1% level, between the maintenance fee measure and the coefficients in columns (1) and (2) of Table 1.6, respectively).[18] This suggests patentees who are most likely to maintain their patents are in the same industries as the patent applicants who sped up their prosecution the most in response to TRIPS. We find similar strong correlations when we compare our measure with patent renewal rates from other years.

[18]This result is not driven by any one particular outlier, as the negative correlation holds at the 1% level when choosing any 36 of the 37 subcategories.

Table 1.6: Post-TRIPS Coefficient and % Maintained Across Subcategories: Chemicals, Computers & Communications, and Drugs & Medical

Subcategory	Post-TRIPS 1	Post-TRIPS 2	% Maintain
Chemicals			
Agric., Food, Textiles	-2.00 (2.72)	-0.39 (3.17)	36.65
Coating	1.07 (1.57)	3.57^b (1.79)	47.83
Gas	0.92 (3.07)	2.06 (3.59)	44.37
Organic Compounds	-4.86^a(1.35)	-1.43 (1.61)	44.30
Resins	-4.99^a(1.21)	-1.73 (1.49)	50.63
Chem.-Misc.	-2.40^a(0.68)	-1.18 (0.79)	45.93
Comp. and Comm.			
Communications	-4.87^a(0.77)	-4.09^a(0.93)	60.68
Comp. Hard. & Software	-5.03^a(0.81)	-4.74^a(0.99)	60.72
Computer Peripherals	-4.17^a(1.37)	-4.03^b (1.73)	62.70
Information Storage	-5.22^a(1.03)	-6.53^a(1.27)	60.41
Elec. Bus. Meth. & Soft.	-6.39^a(2.23)	-6.59^b (2.75)	64.57
Drugs and Medical			
Drugs	-8.55^a(1.00)	-2.56^b (1.23)	48.75
Surgery & Med Inst.	1.65 (1.02)	5.17^a(1.16)	57.81
Genetics	-25.75^a(5.11)	-18.0^a(6.69)	61.66
Drugs and Med.-Misc.	-1.45 (2.26)	0.53 (2.62)	53.47

Post-TRIPS 1 & 2: Robust SEs in ( ). Post-TRIPS 1: 1994-1996; Post TRIPS 2: Excludes 12/8/94-6/7/95. % Maintain = % maintained through 11 1/2 years for patents filed in 1994-1996, by subcategory. ^{a}P$<$0.01,^{b}P$<$0.05,^{c}P$<$0.1.

Table 1.7: Post-TRIPS Coefficient and % Maintained Across Subcategories: Mechanical, Electrical, and Others

Subcategory	Post-TRIPS 1	Post-TRIPS 2	% Maintain
Electrical			
Electrical Devices	-1.18 (1.01)	-0.53 (1.14)	51.65
Electrical Lighting	-1.40 (1.43)	0.30 (1.63)	41.53
Measuring and Testing	-1.99^{c} (1.11)	-2.62^{b} (1.30)	47.31
Nuclear & X-rays	0.03 (1.46)	0.54 (1.64)	51.82
Power Systems	-2.51^{b} (0.98)	-2.34^{b} (1.14)	53.02
Semiconductor Devices	-7.36^{a}(0.98)	-6.08^{a} (1.18)	63.84
Elec.-Misc.	-0.74 (1.17)	0.55 (1.37)	57.00
Mechanical			
Mat. Proc. & Handling	1.11 (0.95)	1.60 (1.06)	41.37
Metal Working	0.80 (1.14)	1.10 (1.28)	45.91
Motors & Engines + Parts	-0.28 (0.92)	-0.42 (1.07)	45.57
Optics	-2.20 (1.55)	-1.86 (1.87)	50.12
Transportation	-1.23 (1.07)	-0.89 (1.23)	36.18
Mech.-Misc.	1.39 (0.86)	2.07^{b} (0.98)	41.09
Others			
Agric., Husbandry, Food	-0.61 (1.45)	-0.79 (1.64)	35.64
Amusement Devices	4.50^{a}(1.67)	5.87^{a}(1.89)	28.64
Apparel & Textile	-1.90 (1.66)	-1.36 (1.91)	30.39
Earth Working & Wells	1.50 (1.83)	2.60 (2.08)	49.53
Furniture, House Fixtures	-1.02 (1.29)	-0.01 (1.45)	27.94
Heating	-1.00 (2.02)	-1.36 (2.46)	40.48
Pipes & Joints	-0.49 (2.27)	-0.93 (2.53)	47.82
Receptacles	-1.10 (1.31)	-0.62 (1.45)	32.98
Other-Misc.	-0.91 (0.68)	0.02 (0.79)	40.88

Post-TRIPS 1 & 2: Robust SEs in ( ). Post-TRIPS 1: 1994-1996; Post TRIPS 2: Excludes 12/8/94-6/7/95. % Maintain = % maintained through 11 1/2 years for patents filed in 1994-1996, by subcategory. ^{a}P$<$0.01,^{b}P$<$0.05,^{c}P$<$0.1.

Table 1.8: Difference-in-Difference with Maintenance Fee Measure

	(1)	(2)	(3)	(4)	(5)	(6)
Post-TRIPS *Maintain	-13.85^a (1.63)	-11.22^a (1.90)	-11.88^a (4.51)	-12.20^a (3.69)	-6.35^b (3.07)	-6.86^b (3.28)
Post-TRIPS	3.26^a (0.81)	3.81^a (0.93)	1.90 (2.42)	-0.74 (2.11)	0.93 (1.60)	1.91 (1.69)
Maintain	28.95^a (1.20)	26.45^a (1.55)	15.23** (7.66)	16.83^a (6.40)	15.28^a (4.18)	15.34^a (4.53)
Time			0.016^a (0.003)	-0.003 (0.002)		
Time *Post-TRIPS			-0.012^a (0.003)	0.007^a (0.002)		
# OAs	106.1^a (0.23)	105.2^a (0.25)	104.8^a (0.90)	104.1^a (0.76)	104.2^a (0.96)	103.4^a (0.77)
# Restricts.			3.95^a (1.09)	4.27^a (1.22)	1.34^a (0.48)	1.84^a (0.52)
Small Entity			6.16^a (0.82)	6.05^a (0.85)	0.96 (0.98)	0.95 (1.18)
Pro Se			-1.11 (0.73)	-1.05 (0.76)	0.58 (0.70)	0.58 (0.84)
Remainder Pendency			0.02^a (0.003)	0.02^a (0.002)	0.01^a (0.001)	0.01^a (0.001)
U.S. Patentee			-19.33^a (0.96)	-20.04^a (0.84)	-0.23 (6.29)	-0.28 (7.50)
Constant	61.72^a (0.61)	61.87^a (0.77)	66.36^a (4.14)	71.14^a (3.30)	65.38^a (3.93)	65.33^a (4.44)
Years	1994-96	No inner window	1994-96	No inner window	1994-96	No inner window
Fixed Eff.	No	No	No	No	Yes	Yes
Adj. R-sq.	0.58	0.58	0.61	0.61	0.68	0.68
Obs.	328,766	258,076	267,372	209,033	267,372	209,033

Appl. delay is dep. variable. Clust. SEs at pat. assignee level in ( ) in all cols. but (1), (2) (robust SEs). Maintain = % maintained over 11 1/2 years for patents filed 1994-96 w/in pat. class, excl. classes w/$\leq$ 100 elements (352 total classes). Cols. (2), (4), (6) excl. appls. filed 12/8/94-6/7/95. Fixed eff. for pat. assignees. ^{a}P$<$0.01,^{b}P$<$0.05.

We also perform a difference-in-difference analysis to test the effect of our previous measure of patent value — the proportion of patentees within an industry who pay all three patent maintenance fees — on applicant delay post-TRIPS. We calculate our maintenance fee measure at the patent class level. The results are shown in Table 1.8, and we can see a strong and statistically significant effect for the interaction between our measure and a post-TRIPS dummy variable. This result remains robust under a wide range of specifications.

The correlation between our measure of patent delay and the renewal rate within a particular industry supports our model of patenting behavior. Namely, patent classes that are most likely to be maintained — and hence, have the highest profit remaining near the end of term — are the same patent classes whose applicants responded the most to TRIPS. In other words, the correlation between these measures suggests that post-TRIPS, applicants who had more profits to lose at the end of their patent term were the ones most likely to speed up patent prosecution. This is what our patent model had predicted.

1.6 Conclusion

In this paper, we use a natural experiment to measure applicants' sensitivity to changes in patent term. In particular, we examine how applicants in different industries differentially sped up patent prosecution in response to the TRIPS agreement, which decreased patent length one day for every day the patent was prosecuted.

To calculate this measure, we construct a novel dataset based on the transaction histories of 331,162 issued patents filed between 1994 and 1996. We apportion prosecution time between the PTO and applicant, and create the first measure of applicant prosecution time. We find that applicants in different industries varied greatly in their response to the law. Our results generally support previous theory, with the exception that computer and software products appear more sensitive to patent term than previously appreciated.

These results are also robust across a number of different specifications. We deal with potential selection around the date TRIPS went into effect by excluding various inner windows around that date. We also exploit any selection itself as an independent cross-industry measure of patent term sensitivity, and find it to be highly correlated with our applicant delay measure.

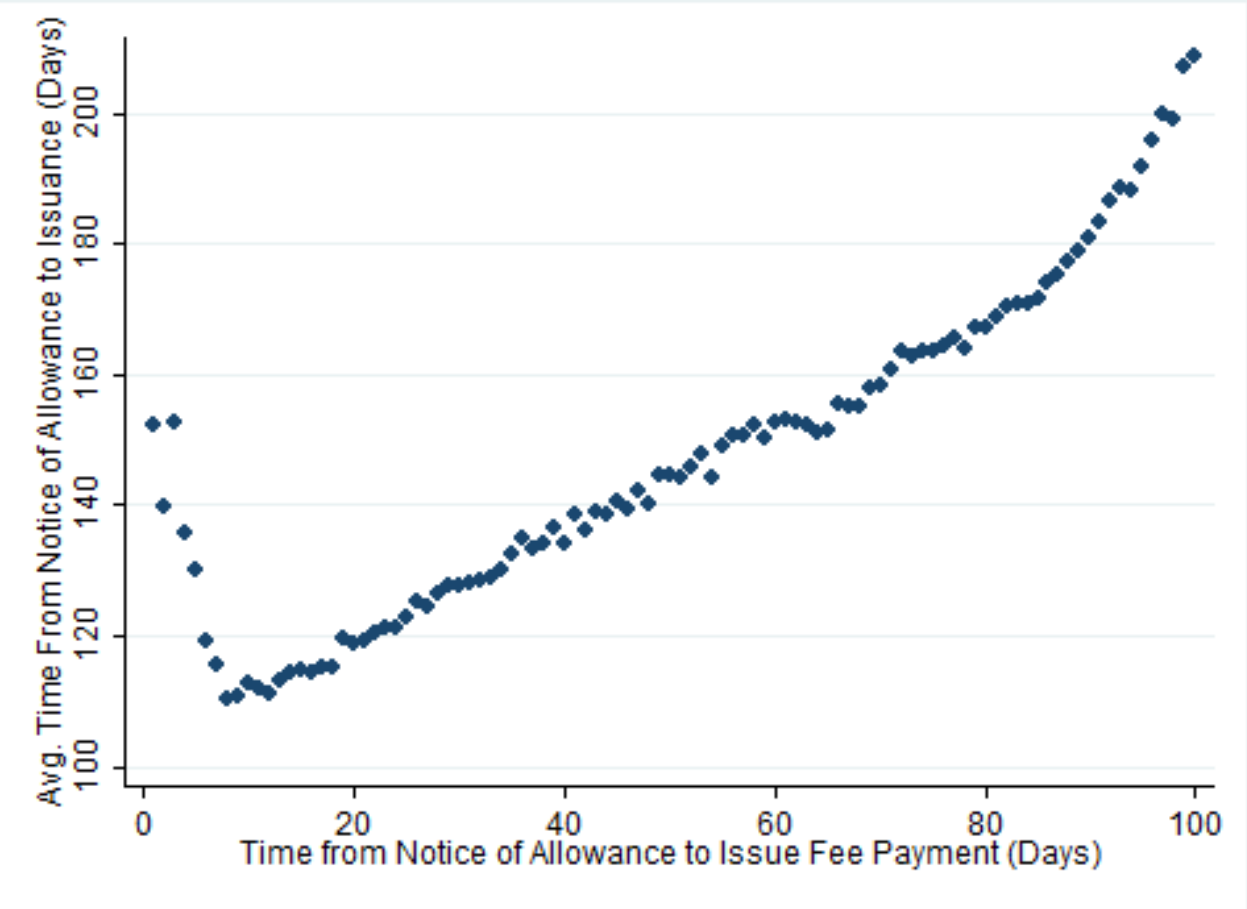
We then present a model predicting that industries with higher profits toward the end of patent term are the industries in which applicants are more likely to speed up patent prosecution in response to TRIPS. To test this model, we develop an independent measure of late-term patent profits based on the proportion of patentees within industry categories who pay maintenance fees through $\mathbf{11\frac{1}{2}}$ years. Because a patentee would likely maintain only those patents whose discounted expected value exceeds the maintenance fee threshold, this measure depends on the profits a patentee expects to receive in the tail end of patent term.

We find our maintenance fee measure is significantly negatively correlated with our applicant delay measure, suggesting that industries in which applicants sped up prosecution the most are the same industries in which patentees were most likely to maintain their patents. This correlation matches our model's prediction and gives us more confidence we have taken a strong step toward credibly assessing the relative importance of patent term across industries.

Appendix

A Additional Figures and Tables

Figure 1A.1: Relationship Between Issue Fee Payment Date and Average Patent Issuance Date



Issue fee payment date and average patent issuance date both measured from mailing date of notice of allowance. Based on data from 472,010 patents issued in 1994-1998. Graph excludes 1.31% (6,183 of 472,010 patents), for which time between mailing of notice of allowance and issue fee payment exceeded 100 days. Such applications likely correspond to extraordinary cases, such as where applicants failed to pay their issue fee in a timely manner and had to petition the PTO for late payment. Regardless, a similar linear relationship appears when graphing dates corresponding to these excluded patents.

Table 1A.1: Applicant Delay by Category with Time Trends

	Chem.	Comp.	Drugs	Elec.	Mech.	Others
	(1)	(2)	(3)	(4)	(5)	(6)
Post-TRIPS	-3.30	-7.03^a	-2.72	-2.27	-1.94	-4.00^b
	(2.04)	(1.83)	(2.96)	(1.71)	(1.71)	(1.77)
time*Post-TRIPS	0.001	0.011^b	-0.011	-0.001	0.007	0.006
	(0.005)	(0.005)	(0.008)	(0.005)	(0.004)	(0.005)
time	0.003	-0.011^b	0.018^a	0.002	-0.004	-0.001
	(0.005)	(0.005)	(0.007)	(0.004)	(0.004)	(0.004)
# OAs	105.9^a	104.3^a	117.2^a	103.4^a	102.8^a	101.1^a
	(0.76)	(0.51)	(0.96)	(0.51)	(0.50)	(0.55)
Constant	75.72^a	80.61^a	70.61^a	74.34^a	73.33^a	72.23^a
	(1.11)	(1.09)	(1.52)	(0.90)	(0.82)	(0.87)
R-squared	0.58	0.62	0.56	0.60	0.56	0.56
F-Test	2.40	7.94	4.95	2.58	1.15	2.64
Prob > F	**0.091**	**3.58e-4**	**0.007**	**0.076**	**0.318**	**0.071**
Est. effect on 6/8/95	**-3.30**	**-1.63**	**-8.76**	**-2.85**	**1.73**	**-0.72**
Observations	39,455	50,044	25,746	49,551	46,430	48,749

Dep. var. is appl. delay. Robust SEs in ( ). Years: 1994-96, excl. 12/8/94-6/7/95. F-Test is for Post-TRIPS and time*Post-TRIPS.
^{a}P<0.01,^{b}P<0.05.

Table 1A.2: Patent Classes w/Post-TRIPS Coeff. Significant at 5% Level

#	Patent Class	Post-TRIPS
30	Cutlery	-8.43 (3.50)
49	Movable or removable closures	10.19 (4.98)
65	Glass manufacturing	14.77 (5.89)
83	Cutting	12.86 (5.56)
100	Furnaces	9.24 (4.46)
112	Sewing	12.08 (4.96)
114	Ships	-7.65 (3.60)
117	Single-crystal, oriented-crystal, and epitaxy growth processes	-10.71 (4.36)
132	Toilet	-12.89 (4.78)
141	Fluent material handling w/receiver or receiver coacting means	-10.27 (4.05)
181	Acoustics	-11.62 (5.45)
187	Elevator, industrial lift truck, or stationary lift for vehicle	-14.87 (3.86)
188	Brakes	-11.01 (3.36)
192	Clutches and power-stop control	13.51 (4.12)
198	Conveyors: power-driven	9.40 (3.44)
203	Distillation: processes, separatory	23.52 (8.62)
205	Electrolysis: processes, comps. and methods of prep. comps.	-11.48 (4.26)
239	Fluid sprinkling, spraying, and diffusing	6.30 (2.87)
257	Active solid-state devices	-7.11 (1.58)
331	Oscillators	-11.83 (3.87)
340	Communications: electrical	-5.36 (2.18)
345	Comp. graphics processing & selec. visual display systems	-4.11 (2.10)
349	Liquid crystal cells, elements and systems	-7.90 (2.84)
358	Facsimile and static presentation processing	-13.84 (2.92)
365	Static information storage and retrieval	-7.01 (1.74)
368	Horology: time meas. systems or devices	-12.27 (5.71)

Dep. var. is appl. delay using spec. (1) in Table 1.1. Robust SE in ( ).
Selected from 294 PTO classes with at least 100 issued patents from applications filed in 1994-96.

Table 1A.3: Patent Classes w/Post-TRIPS Coeff. Significant at 5% Level (cont'd)

#	Patent Class	Post-TRIPS
376	Induced nuclear reactions	9.28 (4.49)
379	Telephonic communications	-10.35 (2.55)
382	Image analysis	-6.29 (2.62)
396	Photography	-10.50 (2.05)
422	Chem. appar. & process disinfecting, deodorizing, preserving or sterilizing	-7.60 (2.89)
424, 514	Drug, bio-affecting and body treating compositions	-7.28 (1.22)
429	Chem: Elect. current prod. appar., product, & process	-9.25 (3.09)
435	Chem: Molecular biology and microbiology	-11.88 (1.73)
436	Chem: Analytical and immunological testing	-12.18 (3.97)
438	Semiconductor device mfg.: process	-7.32 (1.45)
452	Butchering	13.06 (5.45)
454	Ventilation	-12.07 (6.10)
473	Games using tangible projectile	5.72 (2.53)
504	Plant protecting and regulating compositions	-8.59 (4.26)
510	Cleaning compositions for solid surfaces, etc.	-16.71 (3.99)
521, 522	Synthetic resins and nat. rubbers	-11.70 (3.66)
536, 564	Organic compounds	-12.16 (3.51)
705	Financial, bus. practice, management, or cost/price determination	-9.38 (4.02)
707	Database and file mgmt. or data structures	-6.34 (3.00)
708	Comp: Arithmetic process. and calculating	-7.28 (3.37)
710	Comp: Input/output	-9.87 (2.77)
711	Comp: Memory	-5.63 (2.80)
712	Process. architec. and instruction process.	-8.62 (2.93)
717	Software dev., installation, and management	-9.89 (4.37)
720	Dynamic optical info. storage or retrieval	-14.11 (5.59)
800	Multicellular living organisms	-25.75 (5.11)

(cont'd) Classes 521, 522; 536, 564; and 424, 514 are all individually significant but are shown together b/c they share the same description. Classes 705, 707, 717 were all listed as part of "Data Processing."

B Proofs

B.1 Proof of Lemma 1

Recall equation 4 on page 45:

$$-\pi(\tau^*_{new} + \tau_{PTO}) = \frac{dC(\tau)}{d\tau} \mid_{\tau=\tau^*_{new}}$$

Since $\pi(\tau^*_{new} + \tau_{PTO}) \geq 0$, we know $\frac{dC(\tau)}{d\tau} \mid_{\tau=\tau^*_{new}} \leq 0$. This implies $\tau^*_{new} \leq \tau_{min}$, since $\frac{dC(\tau)}{d\tau} > 0 \; \forall \tau > \tau_{min}$.

B.2 Proof of Proposition 1

Because of the applicant's revealed preferences, we know under the old patent term regime (in which patent term is fixed at T_{old} and begins when the patent is issued), the following inequality holds:

$$\int_{\tau^*_{old}+\tau_{PTO}}^{\tau^*_{old}+\tau_{PTO}+T_{old}} \pi(t)dt - C(\tau^*_{old}) \geq \int_{\tau^*_{new}+\tau_{PTO}}^{\tau^*_{new}+\tau_{PTO}+T_{old}} \pi(t)dt - C(\tau^*_{new}) \iff$$

$$\int_{\tau^*_{old}+\tau_{PTO}}^{\tau^*_{new}+\tau_{PTO}} \pi(t)dt - \int_{\tau^*_{old}+\tau_{PTO}+T_{old}}^{\tau^*_{new}+\tau_{PTO}+T_{old}} \pi(t)dt \geq C(\tau^*_{old}) - C(\tau^*_{new}) \tag{7}$$

Similarly, under the new patent term regime (in which patent term is calculated from the patent application date and depends on how long the patent is pending), the following inequality holds:

$$\int_{\tau^*_{new}+\tau_{PTO}}^{T_{new}} \pi(t)dt - C(\tau^*_{new}) \geq \int_{\tau^*_{old}+\tau_{PTO}}^{T_{new}} \pi(t)dt - C(\tau^*_{old}) \iff$$

$$\int_{\tau^*_{old}+\tau_{PTO}}^{\tau^*_{new}+\tau_{PTO}} \pi(t)dt \leq C(\tau^*_{old}) - C(\tau^*_{new}) \tag{8}$$

Accordingly, we know:

$$\int_{\tau^*_{old}+\tau_{PTO}}^{\tau^*_{new}+\tau_{PTO}} \pi(t)dt \leq C(\tau^*_{old}) - C(\tau^*_{new}) \leq \int_{\tau^*_{old}+\tau_{PTO}}^{\tau^*_{new}+\tau_{PTO}} \pi(t)dt \; - \int_{\tau^*_{old}+\tau_{PTO}+T_{old}}^{\tau^*_{new}+\tau_{PTO}+T_{old}} \pi(t)dt \tag{9}$$

We know if $\int_{\tau^*_{old}+\tau_{PTO}+T_{old}}^{\tau^*_{new}+\tau_{PTO}+T_{old}} \pi(t)dt < 0$, then $\tau^*_{new} < \tau^*_{old}$. Now suppose $\tau^*_{old} < \tau^*_{new}$. Then $\int_{\tau^*_{old}+\tau_{PTO}+T_{old}}^{\tau^*_{new}+\tau_{PTO}+T_{old}} \pi(t)dt \geq 0$.

But if $\int_{\tau^*_{old}+\tau_{PTO}+T_{old}}^{\tau^*_{new}+\tau_{PTO}+T_{old}} \pi(t)dt > 0$, then we have a contradiction in equation 9.

So $\tau^*_{new} \leq min(\tau^*_{old}, \tau^*_{min})$ if $\int_{\tau^*_{old}+\tau_{PTO}+T_{old}}^{\tau^*_{new}+\tau_{PTO}+T_{old}} \pi(t)dt \neq 0$ (applying Lemma 1).

B.3 Proof of Proposition 2

From equations 7 and 8 on the preceding page, we know if $\int_{\tau^*_{old}+\tau_{PTO}+T_{old}}^{\tau^*_{new}+\tau_{PTO}+T_{old}} \pi(t)dt = 0$, then:

$$\int_{\tau^*_{old}+\tau_{PTO}}^{\tau^*_{new}+\tau_{PTO}} \pi(t)dt \leq C(\tau^*_{old}) - C(\tau^*_{new}) \leq \int_{\tau^*_{old}+\tau_{PTO}}^{\tau^*_{new}+\tau_{PTO}} \pi(t)dt \iff$$

$$\int_{\tau^*_{old}+\tau_{PTO}}^{\tau^*_{new}+\tau_{PTO}} \pi(t)dt = C(\tau^*_{old}) - C(\tau^*_{new})$$

So we know equation 7 is now an equality and hence, the payoff from choosing τ^*_{old} or τ^*_{new} is the same under the old patent term regime. Similarly, equation 8 is now an equality, so the payoff from choosing τ^*_{old} or τ^*_{new} is the same under the new patent term regime. Accordingly, the applicant is indifferent between τ^*_{old} and τ^*_{new}, regardless of the patent term regime.

References

David S. Abrams. Did TRIPS Spur Innovation? An Analysis of Patent Duration and Incentives to Innovate. *University of Pennsylvania Law Review*, 157(6):pp. 1613–1647, 2009. ISSN 00419907. URL `http://www.jstor.org/stable/40380275`.

David S. Abrams, Ufuk Akcigit, and Jillian Popadak. Understanding the Link between Patent Value and Citations: Creative Destruction or Defensive Disruption? April 8, 2013.

John R. Allison, Mark A. Lemley, Kimberly A. Moore, and R. Derek Trunkey. Valuable Patents. *Geo. Lj*, 92:435, 2003.

Charles Anthon. *A Classical Dictionary: Containing An Account Of The Principal Proper Names Mentioned in Ancient Authors, And Intended To Elucidate All The Important Points Connected With The Geography, History, Biography, Mythology, And Fine Arts Of The Greeks And Romans Together With An Account Of Coins, Weights, And Measures, With Tabular Values Of The Same*. Harper & Bros, 1841.

James Bessen. The Value of U.S. Patents by Owner and Patent Characteristics. *Research Policy*, 37(5):932–945, 2008.

James Bessen, Jennifer Ford, and Michael Meurer. The Private and Social Costs of Patent Trolls. *Regulation*, 34(4):26, 2012.

Michele Boldrin and David K Levine. *Against Intellectual Monopoly*. Cambridge University Press Cambridge, 2008.

Eric Budish, Benjamin N. Roin, and Heidi Williams. Do Fixed Patent Terms Distort Innovation? Evidence from Cancer Clinical Trials. Technical report, National Bureau of Economic Research, 2013.

Ian A. Calvert and Michael Sofocleous. Interference Statistics for Fiscal Years 1992 to 1994, 1995.

Dennis Crouch. Timing of the Notice of Allowance, Issue Fee Payment, and Patent Issuance. http://www.patentlyo.com/patent/2010/06/timing-of-....html, June 9, 2010.

Alan Devlin and Neel U. Sukhatme. Self-Realizing Inventions and the Utilitarian Foundation of Patent Law. *William and Mary Law Review*, 51(3):pp. 897–955, 2009.

Joseph Farrell, John Hayes, Carl Shapiro, and Theresa Sullivan. Standard Setting, Patents, and Hold-up. *Antitrust Law Journal*, 74(3):603–670, 2007.

Michael D. Frakes and Melissa Wasserman. Does Agency Funding Affect Decisionmaking?: An Empirical Assessment of the PTO's Granting Patterns. *Vanderbilt Law Review*, 66(1): pp. 67–149, 2013.

Henry Grabowski. Patents, Innovation and Access to New Pharmaceuticals. *Journal of International Economic Law*, 5(4):849–860, 2002. doi: 10.1093/jiel/5.4.849. URL http://jiel.oxfordjournals.org/content/5/4/849.abstract.

Bronwyn H. Hall, Adam B. Jaffe, and Manuel Trajtenberg. The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. Technical report, National Bureau of Economic Research, 2001.

Jean O. Lanjouw, Ariel Pakes, and Jonathan Putnam. How to Count Patents and Value Intellectual Property: The Uses of Patent Renewal and Application Data. *The Journal of Industrial Economics*, 46(4):405–432, 1998. ISSN 1467-6451. doi: 10.1111/1467-6451.00081. URL http://dx.doi.org/10.1111/1467-6451.00081.

Mark A. Lemley. An Empirical Study of the Twenty-Year Patent Term. *AIPLA Quarterly Journal*, 22(3-4):pp. 369–424, 1994.

Richard C. Levin, Alvin K. Klevorick, Richard R. Nelson, Sidney G. Winter, Richard Gilbert, and Zvi Griliches. Appropriating the Returns from Industrial Research and Development. *Brookings Papers on Economic Activity*, 1987(3):pp. 783–831, 1987. ISSN 00072303. URL http://www.jstor.org/stable/2534454.

Kimberly A Moore. Worthless Patents. *George Mason University School of Law Working Papers Series*, page 27, 2005.

Benjamin N. Roin. The Case for Tailoring Patent Awards Based on the Time-to-Market of Inventions. *UCLA Law Review*, 2014.

Mark Schankerman. How Valuable is Patent Protection? Estimates by Technology Field. *The RAND Journal of Economics*, 29(1):pp. 77–107, 1998. ISSN 07416261. URL http://www.jstor.org/stable/2555817.

Carl Shapiro. Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard Setting. In *Innovation Policy and the Economy, Volume 1*, pages 119–150. MIT Press, 2001.

Neel U. Sukhatme. Regulatory Monopoly and Differential Pricing in the Market for Patents. *Washington and Lee Law Review*, 71, 2014.

U.S. Patent & Trademark Office. Receipts and Dispositions by Technology Center for Ex Parte Appeals FY 1998, 1998.

Charles E. Van Horn. Effects of GATT and NAFTA on PTO Practice. http://www.uspto.gov/web/offices/com/doc/uruguay/URPAPER.html, 1995.

R. Polk Wagner. On Software Regulation. *Southern California Law Review*, 78:457, 2006.

Chapter 2

The Effect of Uber on the Employment and Wages of Taxi and Limousine Drivers

2.1 Introduction

How has the emergence of popular ride-sharing companies like Uber and Lyft affected the wages and employment of taxi and limousine drivers? There has been much speculation that the emergence of the so-called sharing economy is eroding labor standards and lowering wages (Reich 2015; Kuttner 2013), but academic evidence remains scant. Conventional theory does not provide clear guidance about the effect of the entry of transportation network companies (TNCs) such as Uber on the highly regulated labor market for taxi and limousine drivers (Frankena and Pautler 1986). If regulations such as licensing or caps on lease rates raise barriers to entry into driving occupations and provide other protections for drivers, one might expect that drivers would be well positioned to extract rents from management (Acemoglu 2001; Kleiner 2000). On one hand, it follows from the standard rent-sharing model that competition from TNCs would reduce total rents to incumbent firms, thereby lowering wages and rates of employment in the industry (Blanchflower et al. 1992). On the other hand, if regulations are nonbinding or if capital holders have access to a large pool of additional workers, then drivers' wages should equate to their marginal product. Under that competitive framework, the entry of TNCs with better technology and less regulatory burdens (Cramer and Krueger 2016) should lead to an outward shift in demand for rides,

with the rates of pay of workers in the industry as a whole increasing accordingly and with the magnitude depending on the elasticity of the supply of drivers (Lanzona and Evenson 1997). This paper updates the first nationwide evidence on the effects of popular mobile ride-hailing service Uber on the wages and employment of taxi and limousine drivers.[1]

In contrast to studies that have focused on Uber's experience in a single market (Schneiderman 2015; Weise 2015), I analyze the experiences of taxi and limousine drivers in the 70 largest Core-Based Statistical Areas in the United States from 2008 through 2016. I begin by presenting a straightforward difference-in-difference estimation procedure by using the varied timing and size of Uber's entry into different regions to evaluate Uber's effect on driving occupations, and I then extend the estimation strategy to include fixed-effects models with flexible preexisting trends to control for city and time effects.

In contrast to several claims in the media, I find no evidence to support the claim that Uber has significantly driven down the wages of taxi drivers and chauffeurs writ large. This is in stark contrast to past literature on deregulation (Card 1998; Hirsch 1993; Rose 1987), which focuses on industries with strong unions to extract rents created by regulators. Even on an hours-adjusted basis, the launch of Uber has resulted in an increase in the overall employment of taxi drivers and chauffeurs. That finding is important in its own right, because in the second half of 2016, almost 600,000 workers reported their main or secondary occupation to be taxi driver or chauffeur. It is perhaps more valuable, however, for what it implies about the competitiveness of the labor market and the likely effects of significant industry disruption – caused by technological innovation – on the compensation of existing workers in the industry.

In addition to the empirical results, a few aspects of the analysis are of interest for broader research. To my knowledge, this is the first paper that uses Uber's administra-

[1] Although there are other competitors in the marketplace, Uber is the clear leader; and focusing on one company aids in identification strategy.

tive data on the number of active drivers by city and month for 18 of the 21 largest US metropolitan statistical areas (MSAs) – through the end of 2014 – to measure Uber's differential impact on the taxi industry across regions.[2] I first present estimates for this relatively small sample of cities, wherein the data are most accurately measured. In this sample, I use difference-in-difference estimation and more-sophisticated econometric estimation and find no evidence of negative effects – based on Uber's market penetration – on the wages of taxi and limousine drivers in that area. Next, I extend the sample of cities by creating an original data set of the timing of Uber's launch of different services in 50 additional cities provided to me by Uber.

Finally, I derive a new measure of Uber's penetration in these cities through the end of 2016 by creating a city-level data set of the volume of Uber-related Google searches. The application of Google search data, derived from use of the Google Trends tool, as a proxy for a variable that is otherwise not available was originated by Choi and Varian (2012) and used by Stephens-Davidowitz (2014). In the current application, I can accurately assess whether Google search activity is capturing variation in the actual number of Uber drivers for the 18 MSAs provided in Hall and Krueger (2015). I use that sample to scale Google searches in terms of the likely number of Uber drivers actively providing services in a city, to allow for a two-sample two-stage least squares (TS2SLS) estimation procedure. My preferred estimates come from this largest sample.

My analysis suggests that competition from Uber has caused large capital losses among medallion holders and taxi and limousine businesses more generally but not among the drivers those businesses employed or contracted with. Although representative data are not available, it appears that in many metropolitan areas, the value of a taxi medallion has fallen by at least 30% from its peak around the beginning of 2013 – before Uber was a

[2]Administrative data on the number of Uber drivers in a city each month are drawn from Hall and Krueger (2015). Data for Detroit, Tampa, and St. Louis were not included; they are the 14th-, 18th-, and 19th-largest MSAs, respectively.

major player. Overall, my results provide little support for the hypothesis that the wages of workers in taxi and limousine services have fallen significantly as a result of Uber's entry. That is inconsistent with rent sharing's being predominant in the industry but is in line with competitive forces determining industry wages.

The remainder of the paper is organized in sections. Section 2 provides background information on Uber and the taxi and limousine industry. Section 3 gives an overview of the related literature on deregulation and the recent literature on the sharing economy. Section 4 offers two simple models of ways Uber might affect the labor market for taxi drivers. Section 5 describes in detail the data that were constructed for this paper and their potential shortcomings. Section 6 presents the empirical results along with several robustness checks. Section 7 consists of concluding comments.

2.2 Background

2.2.1 Background on Uber

Uber was founded in the beginning of 2009 as an app-enabled, ride-hailing platform.[3] Anyone with a smartphone and a credit card who downloads the Uber app can launch the app to request rides from drivers who are also using the platform in the area. Originally started solely as a black-car service in San Francisco, Uber launched its lower-cost UberX service – by far its most popular – in mid-2012. The base fare for customers using UberX is usually set to be 10 to 20% lower than the comparable taxi fare schedule in the city, and drivers can use their everyday cars. The company also offers UberSUV, UberLUX, and UberPOOL in

[3]Techcrunch.com/topic/company/uber/. Many of the sources for this paper are websites as information has changed quickly.

many markets.[4] From very few drivers in mid-2012, the numbers of drivers and riders using the service have increased exponentially, with estimates that there are currently more than 400,000 Uber drivers on the road domestically by the end of 2015.[5][6] Recent data has not been released, but speculation and my predictive model would suggest that the number of US Uber driver-partners stands currently around one million. Uber has been one of the top five most downloaded travel apps in the United States every week since mid-2013.[7] Figure 2.1 shows the pace of growth in number of drivers on the Uber platform through the end of 2014 as well as the continuing increase in traffic for Uber-related Google searches as the company's market penetration deepened. Still a private company, Uber's most recent valuation was $69 billion, and the service is available in more than 450 cities worldwide.[8]

In its current incarnation, at the time of writing, the app shows users (1) the time it would take for a car to reach their location based on different vehicle and service types and (2) a projected fare to their destination. The fares generally follow a standard formula, but when Uber's algorithm determines that more people are accessing the app relative to the number of drivers in the area, prices for some services surge – by multiples of the fare and often with a minimum fare requirement in the past. The increased cost reduces oversubscription when more people are on the system than can be offered rides from drivers on the app in a given amount of time. A pair of case studies compared the night of a sold-out concert at Madison Square Garden with the early hours of New Year's Day, when surge pricing failed briefly in New York. The analysis showed that during a surge, relatively fewer rides are requested and more drivers use the platform than would otherwise be the case (Hall et al. 2015). For this study, surges represent opportunities for drivers to earn more money

[4]http://newsroom.uber.com/johannesburg/2014/10/uber-pro-tips/, https://www.uber.com/, http://newsroom.uber.com/announcing-uberpool/, for, respectively, "those times when you need a bit more space"; "The finest cars"; and, "share a ride – and split the cost."

[5]https://newsroom.uber.com/2015/06/5-years-travis-kalanick/.

[6]http://www.1776.vc/1776-and-uber-live-uber-and-the-american-worker-with-ezra-klein-david-plouffe/.

[7]App Annie Index 2017.

[8]https://www.bloomberg.com/gadfly/articles/2017-03-16/uber-needs-to-get-real-about-that-69-billion-price-tag, accessed 10/8/2015.

on an hourly basis, ceteris paribus, and to follow predictable patterns. Demand for Uber is reported to peak around the time of bar closings, for example.[9] Recent investigations of the Portland, Oregon, taxi market show that Uber takes a larger share of for-hire trips on weekend nights.[10] Payments for Uber rides are processed electronically with the user's previously entered credit card information, and tipping while discouraged in the past has been allowed recently.[11] Uber retains a percentage of the fare – the so-called take rate – which is generally 20 to 30%, although it varies by city.[12]

Drivers must qualify to utilize the Uber app – typically, by passing a background check and undergoing a review of their driving record; by submitting documentation of insurance, registration, and a valid driver's license; by successfully completing a city-knowledge test; and by driving a vehicle that meets a quality inspection and is less than a certain number of years old.[13] They must also maintain a satisfactory rating and accept a sufficient number of rides that they are offered while logged in to the app. Data on the number of applications rejected or drivers suspended is not publicly available, but anecdotal evidence suggests few drivers are kicked off the platform once admitted.[14]

Little information has been released publicly about the backgrounds and work habits of Uber drivers. A Benenson Strategy Group survey of 601 Uber drivers in 20 markets found that drivers on the platform were younger and better educated than were taxi drivers and chauffeurs in government surveys covering the same period. Drivers are free to choose the number of hours they work on the platform – an advantage Uber often points out. Although

[9] https://newsroom.uber.com/wp-content/uploads/2015/01/UberMADD-Report.pdf.

[10] http://www.bloomberg.com/news/articles/2015-10-19/these-four-charts-show-how-uber-steals-turf-from-cabs-and-spreads-throughout-a-city. This study also included rides given by Lyft, Uber's largest competitor in the US market.

[11] http://www.businessinsider.com/uber-tipping-policy-2014-1 and http://www.marketwatch.com/story/this-is-exactly-how-much-you-should-tip-your-uber-driver-2017-04-18.

[12] http://www.wsj.com/articles/uber-tests-30-fee-its-highest-yet-143198912.

[13] Hall and Krueger 2015.

[14] http://www.businessinsider.com/leaked-charts-show-how-ubers-driver-rating-system-works-2015-2?r=UK&IR=T.

the vast majority (81%) of taxi drivers and chauffeurs report driving for 35 hours or more per week from 2012 to 2013, according to data from the American Community Survey, the Benenson survey found that 44% of Uber drivers were on the road for less than 12 hours a week in 2014.[15] That flexibility is one of the features that has led to an inflow of drivers from other industries.[16] Although the industry from which most Uber drivers originated is transportation, this accounted for only 19% of total Uber drivers. A slight majority of those who drove on the platform said they had never previously worked as drivers. In addition to driving part-time with Uber, 62% of drivers at the end of 2014 had other jobs, either full-time or part-time. Data on longer-term retention rates of Uber drivers are not yet available, but early evidence shows that drivers' usage of the app can be transitory.[17]

2.2.2 Background on the Taxi and Limousine Industry

The US taxi and limousine industry is subject to idiosyncratic regulations at both the state and local levels that has been a prototypical example of government inefficiency (Coffman 1977; Tullock 1975). In most large cities, drivers must be granted the right to carry passengers in taxis at predetermined fare schedules, which are usually conveyed in the form of a vehicle license called a medallion, which is generally leased to a driver by a taxi and limousine company. The supply of medallions is typically determined by the municipality. The high prices at which medallions were bought and sold prior to Uber's introduction in many major metropolitan areas attest to the current and expected profits that accrue to the medallion holder. In New York, for example, the cost of a minifleet medallion rose to $1,115,000 in

[15] http://www.bsgco.com/insights/uber-the-driver-roadmap.

[16] "To have more flexibility in my schedule and balance my work with my life and family" (85%) was the third-most-common reason listed by the Berenson Strategy Group when Uber drivers were asked why they drove with Uber. Reasons one and two were "to earn more income to better support myself or my family" and "to be my own boss and set my own schedule," respectively (Hall and Krueger 2015).

[17] According to Hall and Krueger (2015), 11% discontinue, defined as becoming inactive on the system for six months' service in the month after starting; 30% discontinue after half a year; and just less than half discontinue after a year. Data cover those who started in the first half of 2013.

2013 from $440,000 in 2004, outpacing the nominal returns of the stock market by orders of magnitude since the end of World War II (NYC Taxi and Limousine Commission 2014).[18] Likewise, in Philadelphia, the price of a medallion rose from $65,000 in 2005 to $475,000 in 2014, but recent auctions have suggested that prices may have fallen back to their previous levels.[19] Many areas also require limousine services to have specially licensed drivers and to pay fees to carry passengers. Data on area-specific limousine requirements are not well publicized.

In addition to fare regulation, the amount medallion holders can charge drivers to lease or rent cars or medallions is limited in certain jurisdictions. In spite of those regulations and the large returns to medallions in the pre-Uber period, government data give only weak evidence of rent sharing. In May 2016, the median hourly wage was $11.68 for taxi drivers and chauffeurs, whereas it was $12.20 for those in other transit and ground passenger transportation in the Occupational Employment Statistics (Bureau of Labor Statistics 2014). Similarly, the Bureau of Labor Statistics' *Occupational Outlook Handbook* reported that in 2012, taxi drivers' and chauffeurs' median annual pay was $22,820 – lower than that of five other occupations reported as having the most-similar job duties and education requirements.[20] The hourly wages of taxi and limousine drivers are about 20% lower than the wages of workers in other industries who have the same education level, age, gender, race, ethnicity, immigration status, and marital status in the Current Population Survey (CPS).[21] Chauffeurs earn a larger share of their wages via tips, a fact that might not get captured in government surveys and that perhaps explains the apparent wage penalty for being taxi drivers

[18] http://www.nyc.gov/html/tlc/downloads/pdf/2014_taxicab_fact_book.pdf.

[19] http://www.newsworks.org/index.php/local/business-a-economy/74302.

[20] The 2012 median pay rates were: bus drivers, $29,550; railroad occupations, $52,400; delivery truck drivers and driver/sales workers, $27,530; automotive service technicians and mechanics, $36,610; and water transportation occupations, $48,980; http://www.bls.gov/ooh/transportation-and-material-moving/taxi-drivers-and-chauffeurs.htm#tab-7.

[21] This comes from a Mincer regression of log hourly wage on four education dummies, experience, experience squared, race, Hispanic ethnicity, marital status, and gender. As will be discussed later, not all taxi drivers and chauffeurs will be included in this regression.

or chauffeurs. A 1997 Bureau of Labor Statistics investigation of tipped workers in New Orleans found that workers in the amusement, passenger transportation, and hotel category had the highest average hourly tip rate – at $7.85 (Simpson 1997).[22] Most employment in the taxi and limousine industry is at will, and short-term contracts predominate. Since 2008, 8% of taxi drivers and chauffeurs reported by way of the CPS that they were union members, with an additional 1% saying they were covered by union agreements. In the same survey, three-quarters of taxi drivers and chauffeurs reported receiving no money in employer contributions for health insurance.

In spite of the lack of evidence that taxi and limousine drivers are relatively well compensated, it has been documented that many drivers vehemently oppose the spread of Uber and other services both in the United States and abroad.[23] The high, fixed costs of job switching and the status quo bias may help explain those drivers' positions. Data on worker flows into the non-TNC sector are not generally available, but in Los Angeles – one of Uber's biggest markets, for example – monthly applications to become licensed taxi drivers were down 66% from before Uber's entry (Los Angeles Department of Transportation 2015).[24] Legal efforts to thwart Uber are ongoing in numerous jurisdictions.

2.3 Related Literature

The literature most closely related to Uber's disruption of the taxi and limousine industry analyzes what happened to wages and employment after the nationwide deregulation of the trucking and airline industries in the 1970s and 1980s. Interstate trucking and air travel had

[22] http://www.bls.gov/opub/mlr/cwc/tips-and-excluded-workers-the-new-orleans-test.pdf.

[23] Perhaps the most widely reported protests were in Paris: http://www.nytimes.com/2015/06/26/business/international/u
protests-france.html. Riots in New York were, according to media reports, led by a union representing New York taxi drivers; http://abc7ny.com/traffic/taxi-drivers-protest-halt-of-uber-cap-in-nyc;-cuomo-says-service-will-go-state-wide/875555/.

[24] Private communication with Commissioner Eric Spiegelman, president, Los Angeles Taxicab Commission.

been heavily regulated, with fares stipulated by government agencies after the end of World War II (Dempsey 2012). In 1978–79, the Interstate Commerce Commission began to pass reforms of the trucking industry, effectively eliminating most price and entry restrictions in the industry by way of the Motor Carrier Act of 1980 (Harper 1980). Deregulation allowed many nonunion companies – previously not permitted to carry cargo – to enter the trucking industry and hire new workers at lower wage schedules. Scholars have investigated the changes in union wage premiums as a proxy for the degree of rent sharing in the industry (Peoples 1998). The first estimates using CPS and labor union data suggested labor was capturing as much as 75% of total rents in the trucking industry prior to reform, such that a representative union driver's annual wages were $3,800, or 14%, higher than they would have been otherwise (Rose 1987). Using a longitudinal data set that included job switchers, Hirsch (1993) found smaller but still negative wage effects of deregulation.

Card (1998) analyzed the wage dynamics of the air transport industry after passage of the 1978 Airline Deregulation Act, which liberalized fare setting and expedited requests for new routes. Car found that deregulations' effects on airline workers caused workers' relative earnings to fall by around 10%, but the union-to-nonunion wage differential remained relatively large. Card argued that it could be either because wages had been relatively competitive before deregulation or because there were still rents to be shared after deregulation.

Whether Uber's effect would be seen in this type of analysis is not clear. As mentioned earlier, the unionization rate in the taxi industry is low, and therefore the mechanism by which rents were won in the trucking and airline industries is not feasible. It is also unclear whether the same data sources used in those analyses would be as reliable for taxi and limousine driver occupations. Many studies have found that in government surveys, people underreport to the Internal Revenue Service their employment and their earnings from self-employment and part-time employment (Hurst et al. 2014). If that's the case for Uber drivers

and many of the new participants in the gig economy, there are important implications for policy makers about what data sets to use to inform their decisions.

Unlike the literature covering many other industries, the theoretical taxi and limousine industry literature is mixed when it comes to the need for strong government regulation of taxis and limousines. Manski and Wright (1967) derived a model that argued that because demand is a function of supply – more taxis on the road lead to lower wait times and to more rides demanded – the system is not stable unless there are regulated prices. On one hand, the need for some regulation of fares, even with free entry, was also suggested by Abe and Brush (1976), De Vany (1975), and Douglas (1972) in order to prevent a situation wherein an individual customer has to bargain with the one taxi driver he stops on the street, thus allowing the taxi driver to assert some degree of monopoly power. Williams (1980), on the other hand, argued for complete deregulation. More recently, Cairns and Liston-Heyes (1996) suggested that "price regulation is essential, and entry regulation may be useful" to the functioning of the market as well.

There have been attempts at deregulation of the taxi industry both in the United States and abroad, with varying outcomes. Barrett (2010) reports on Ireland's positive experience with deregulation after an Irish High Court decision in 2000 that resulted in a quintupling of the number of taxis, with increased output and reduced waiting times. Information on taxi wages in Ireland is hard to compile, but Goodbody Economic Consultants (2009) suggested that wages declined slightly after deregulation. The value of licenses fell dramatically from €114,000 prior to the ruling. Teal and Berglund (1987) summarized the relatively disappointing experience with taxi deregulation in the United States in the 1970s and 1980s.[25] The authors reported that in all of the cities examined, taxi prices had not fallen in real terms, nor had service improved notably. There was no analysis of wage pass-through to individual drivers in the different periods covered by those papers. All of the cities studied

[25]In the 1970s and 1980s, San Diego, Seattle, Sacramento, Kansas City, Oakland (CA), Fresno (CA), Charlotte (NC), Portland (OR), Tucson, and Phoenix all experimented with deregulation.

wound up reversing course on deregulation. Baanders and Canoy (2010) described similar negative experience in the Netherlands, where nationwide deregulation in 2000 resulted in higher taxi fares and worse service. Current TNCs, unlike entrants following earlier complete deregulations, inform consumers of their fare schedules ahead of time, thereby eliminating possible haggling – especially in a street-hail situation, which many saw as destructive to the system.

Other, related studies have focused on the taxi industry to estimate labor supply decisions because taxi driver is one of the few occupations in which workers can choose work hours flexibly and in which the wage forgone by taking leisure is known. Early work by Camerer et al. (1997) suggested that taxi drivers seemed to be target earners and might not be completely rational in their supply decisions, but more-recent work by Farber (2015, 2008) found that in general, drivers don't appear to be target earners and have labor supply elasticities in line with the literature. Ashenfelter et al. (2010) found that the driving schedules of owner-drivers in New York declined after fare increases, suggesting that income effects predominated. Overall, analysis attests to the fluidity of market-time supplied drivers.

A body of literature measuring Uber's effects on various aspects of life has begun to develop. In one of the earliest papers, Greenwood and Wattal (2015) attempted to estimate the effect of the introduction of different Uber services in California counties on alcohol-related vehicular fatalities. The authors found no effect on fatalities after the introduction of UberBlack – the luxury version of UberX – but they found a significant decrease in deaths in counties once UberX was introduced. Edelman and Geradin (2015) analyzed the efficiencies of the software platform embedded in Uber, which they argued are substantial and should be allowed – but regulated similarly to regulated, traditional cab companies. Using New York taxicab data, Frechette et al. (2015) showed that a more efficient matching algorithm – similar to Uber's – has the potential to increase cab drivers' earnings by 2.5% – enough to offset possible losses from a slightly increased number of for-hire vehicles on

the road. Deriving a more complete spatial matching equilibrium, Bucholz (2016) shows that an optimixed directed matching technology would increase welfare in New York City by $2.4 billion per year. Cohen et al. (2016) use discontinuities in how surge prices are rounded to estimate consumer surplus and estimate it to be $6.8 billion for the U.S. in 2016. Berger et al. (2017) builds upon an earlier version of this paper to show that there have been no disemployment effects of Uber. An investigation using ACS data and truck drivers as a control group argues that wages have risen for self-employed drivers but fallen for wage earners. This paper prefers wage estimates from the QCEW which gives a fuller picture of wages of taxi employees and finds little evidence that their wages have fallen.

2.4 Model

In this section, I describe two basic models of wage determination and how their predictions for wage changes differ following the entrance of Uber into a city: the first is a model of a market wherein medallions and cars are held by owners and leased to drivers through competitive bidding; the second is a market model wherein a lease cap or barriers to entry for drivers cause rent sharing between medallion holders and taxi drivers.[26] Both cases consider a taxi market that consists of cars and medallion owners and workers who meet daily. Each car owner can lease a car to a worker for a daily price, represented by c. If a driver chooses to rent a car, total earnings for that car that during the shift are represented by R.[27] A driver who does not choose to drive can earn per shift an outside wage net of job-switching costs, w_0. Workers' utility is solely a function of income, and they do not

[26]For simplicity, I abstract away from the considerations of optimal regulatory or matching algorithms that might follow Uber's development – as explored by Frechette et al. (2015) – and focus on how an individual driver sees the market. This is more similar to Ashenfelter et al. (2010), who assumed an individual driver receives the market wage regardless of hours worked.

[27]I abstract away from the hours the decisions of drivers.

consume leisure. I also assume that it is costless to owners if their cars sit idle. There is a cost of capital to car owners associated with depreciation, represented by D.

2.4.1 Competitive market with unilateral pricing of lease rates, c in the market

If the owner leases a car, the driver earns:

$$Earnings = R - c \tag{10}$$

Drivers drive as long as their earnings from driving exceed their outside opportunity wages.

The car owner can solve the driver's problem and knows that the driver will drive as long as $R - c \geq w_0$. That is, the driver drives as long as that driver earns more from hours driving than would have been earned by spending the same amount of time at that driver's second-best job. The highest rate the owner can charge is thus $c = R - w_0$, at which point the driver is indifferent to renting. Owners will choose to rent at that rate if c exceeds D, the cost of depreciation that the owner bears solely. Thus, when a rental agreement is reached, owners take all profits. In that situation, if revenue function R changes, the profits that accrue to the workers will be the same – namely, zero. The owner will simply resolve the driver's problem and adjust c accordingly. Alternatively, if the outside wage w_0 increases, then the earnings of workers who drive taxis will rise.

2.4.2 Rent-sharing case, with Nash bargaining over profits

The taxi and limousine industry is subject to many regulations covering fares as well as lease rates in many areas (Williams 1980). The need for licensing and regulation could lead to rent sharing, or quasi-rents for workers, a situation I generalize with classic Nash bargaining

(Kleiner 2000; Abowd and Kramarz 1992). Recall from earlier mention that car owners' total earnings consist of the cost charged for leasing the cab, c, net of depreciation costs on their capital, D. Combining that with what cab drivers earn in (1), total profits in the industry are simply:

$$\Pi = R - D - w_0 \tag{11}$$

where c, which is a pure transfer, does not appear here. Profits are total earnings in the industry minus the cost of capital, D, and the opportunity cost of time, w_0. In the generalized Nash bargaining equilibrium (Svejnar 1980), each side has relative bargaining power, which I define as γ for labor and $(1-\gamma)$ for capital.

The generalized Nash solution is the value of c that maximizes the value function, V:

$$V = (R - c - w_0)\gamma(c - D)(1-\gamma) \tag{12}$$

where $(R-c-w_0)$ is the surplus to labor and $(c-D)$ is the surplus to capital. The solution to this maximization problem takes the familiar form leaving an expression for c, the earnings for capital holders,

$$c = (1-\gamma)(R - w_0) + \gamma D \tag{13}$$

Using (1), the total driver net earnings, Πd are:

$$\Pi d = \gamma(R - w_0 - D) \tag{14}$$

Drivers thus receive γ share of total surplus, with capital receiving $(1-\gamma)$ share. With these expressions, one can investigate how a driver's income should change with a change in R. Unlike before, it is clear now that as R decreases, so too do a driver's earnings and profits. It can also be seen that as the outside option w_0 changes, c moves in the opposite direction. I note still that drivers' wages can fall only until they hit the outside wage, w_0, at which point

drivers would leave the industry. An increase in w_0 could thus offset some of the decline in wages because of less profitability. The following sections empirically test whether taxi and limousine drivers' wages decline in the presence of Uber – as predicted by the theory of rent sharing – or whether wages have remained constant or risen, consistent with a competitive model.

2.5 Data Sources

To assess taxi and limousine drivers' performance in areas with Uber, I use a data set on Uber's rollout and size across cities. As mentioned previously, one of the main sources of data is from Hall and Krueger (2015), who included the most-recent Uber driver information released by the company, including numbers of active drivers by city and month for 18 of the 21 largest US metropolitan statistical areas from the middle of 2012 through the end of 2014 (figure 2.2).[28]

I next extend that sample of cities by creating an original data set of the timing of Uber's launches of different services in 60 additional cities. Birmingham, Buffalo, and Rochester, New York, were the largest cities where Uber had never operated in by the end of 2016. To obtain a measure of Uber's relative popularity for this set of 80 MSAs, for which the number of Uber drivers is unavailable, I create a city-level data set of the volume of Uber-related Google searches in each area by month. Originally suggested by Hal Varian as a way to better predict current economic conditions – including automobile sales, home sales, retail sales, and travel behavior (Choi and Varian 2012) – Google Trends, which gives a measure of relative search intensity across time and location, has become a useful tool, especially for data on private companies and other hard-to-measure variables. Recent users have in-

[28]Uber in the past has been very reluctant to release data, something that has gotten the company in trouble with various regulators, but recent signs suggest more data sharing to come: http://blogs.wsj.com/digits/2015/01/13/uber-offers-trip-data-to-cities-starting-in-boston/.

cluded Stephens-Davidowitz (2014), who used Google Trends frequency for a racial slur as a proxy for animus toward African-Americans across regions. Google Trends data measure the relative frequency of Uber-related searches in an area normalized by total searches in that region. For the purpose of this analysis, that produces a metric that is comparable across cities, assuming only that general internet usage is similar.

I investigated terms that would mirror Uber's penetration into different metropolitan areas along the lines of figure 2.2. The simple term Uber, which had very little search activity nationwide prior to the launch of UberX, as shown in figure 2.1, mirrors the rollout both nationwide and across cities, as seen in figure 2.3, and is perhaps the broadest measure of penetration. Figure 2.3 shows that for cities included in Krueger and Hall's survey, Uber-related searches aligned closely with the number of drivers on the platform. Of note is how Miami, shown in figure 2.4, which was a late adopter of Uber, experienced rapid acceleration in Uber searches in line with the city's increase in drivers once the service launched. In the current application, I can assess whether Google search activity is capturing variation in the actual number of Uber drivers for the 18 MSAs with data released by Uber, and I can use that sample to scale Google searches based on the likely number of Uber drivers actively providing services in a city. Regression analysis in the next section confirms the robustness of the relationship and the validity of Google Trends data as a proxy.

To attempt to measure the performance of the taxi and limousine industry, I use two government data sets: the Current Population Survey and the Quarterly Census of Employment and Wages (QCEW). The most up-to-date data set, as well as the data set from which the US Department of Labor determines the unemployment rate, is the Current Population Survey. The CPS is a survey of 60,000 households in which all residents older than age 15 years are asked whether they were employed and what kind of work they did during the week that includes the 12th of the month. In the CPS's rotation group design, individuals are interviewed for four consecutive months, are then dropped out of the survey for eight

months, and are then interviewed again for four more months. That design makes it possible to examine the job flows of those who reported they were employed in taxi and chauffeur occupations – which is one of the categories – for a 15-month interval.

In the CPS there is a well-known bias toward the underreporting of second jobs, as shown most recently by Alm and Erard (2015); so those who are Uber drivers in addition to another job might not show up in those occupations. Workers are asked for their weekly earnings only when they exit the survey – their fourth and eighth months in the sample. For my main regressions, I collapse the observations on wages, employment, and hours in the taxi industry to city-month cells. In many Core-Based Statistical Areas, there were months for which no taxi wage was recorded. It is unclear how taxi and limousine drivers report their incomes – especially if they are self-employed – but that measurement error in the dependent variable should not bias results as long as reporting practices have not changed post-Uber. The main drawbacks of this data set are lack of sample size and uncertainty about how Uber drivers choose to classify themselves. The presence of Uber drivers in the CPS could affect my results in different ways that are discussed more thoroughly later in relation to the regressions. Due to its wide coverage the CPS is my preferred dataset for estimates of total employment effects.

The QCEW is a census of all companies; it is built on administrative unemployment insurance records; and it gives the employment level and average wage at the industry level for each county and MSA at a monthly frequency. Many taxi drivers who work on a daily or weekly basis will not be captured in the survey because they are not formal employees. Data are included for the taxi and limousine services industry, with the number of industry workers being about half that of the taxi industry nationwide, based on a CPS tabulation of workers who give their industry as being taxi and limousine relative to the QCEW nationwide number. Although the occupation of every person covered by the survey is not observable, in the American Community Survey 85% of people in the industry were classified as taxi drivers

or chauffeurs for their occupation. The other workers are distributed between other, support, repair, and supervisory occupations and the second-most-prevalent role: dispatchers (5%). Only workers in the taxi and limousine industry who are covered by unemployment insurance are included in the survey. Unemployment insurance coverage varies across states, but in all cases, those who drive for Uber solely are not included in this survey. For QCEW results, the person to consider is not a driver who leases a car daily but a limousine or taxi driver who works routinely as a company employee. On occasion the Bureau of Labor Statistics does not release data for certain MSAs in certain periods when the bureau feels that disclosure could identify a company or workers. In spite of this, due to its large coverage and administrative data, the QCEW is my preferred source for wage data.

2.6 Empirical Work

Nationally, the number and total hours of drivers in the CPS (figures 2.5 and 2.6) and the real weekly wage rate from the QCEW are all at record highs. It is likely that Uber is the biggest change in the taxi and limousine market but there could have been other factors nationwide. As mentioned in the previous section, the primary data set contains city-level data on the wages and employment of taxi drivers in MSAs as well as measures of Uber's activity. I measure in many ways the wage and employment effects of Uber's proliferation across cities. First, I use company-released data on the number of drivers in the 18 MSAs where Uber was most active as of the end of 2014, and estimate a model of the form:

$$Y_{jt} = \beta_0 + \beta_1 log(Uber\ drivers_{jt}) + \beta_2 urate_{jt} + \mu_j + \alpha_t + \epsilon_{jt} \tag{15}$$

where Y_{jt} is wages or employment in city j in month t, $log(Uber\ drivers_{jt})$ is the logarithm of the number of active drivers on the Uber platform in city j in month t normalized by the total population, $urate_{jt}$ is the unemployment rate in city j in month t, μ_j is a vector

of city-level fixed effects, and α_t is a vector of month-year fixed effects. Of interest is β_1, the elasticity of wages or employment with respect to a 1% increase in the number of Uber drivers relative to the city's population. The results in table 2.1 show that from the inception of Uber to the end of 2014, when Uber became more popular in terms of drivers, there were no significant changes in wages and employment in the QCEW or CPS. The point estimate for the CPS wage regression is slightly negative but far from significant. The employment estimates both have negative point estimates, which are coincidentally the same, but their magnitudes suggest that as the number of Uber drivers relative to the population increased by 10%, hourly wages would decline by just 0.08%. Many transformations of the main independent variable were tested – including different normalizations, by number of taxi and limousine drivers in a city prior to Uber for example – with no meaningful changes in the results. Different weighting schemes – population and pre-Uber taxi and limousine drivers – affected results minimally. The model is robust for the inclusion of city-specific linear and quadratic time trends. Overall, using Uber administrative data, the model shows little evidence of declining wages and employment where Uber had become more popular as of 2014.

The drawbacks of the regressions in Table 2.1 – the most expansive possible using solely Uber-released driver data – are that the data cover only 18 cities and only through the end of 2014. It is possible that the true effect of Uber will take longer to develop – if there are costs involved in switching jobs, for example, or if Uber has to reach a certain maturity level before its full effect is realized.

To extend the data set to measure both the full set of 80 cities and the extent of Uber's popularity, I use as a proxy for the number of drivers the Google Trends data described earlier. In particular, I use search intensity for Uber-related terms as one of the main controls. As suggested by figures 2.3 and 2.4, Google Trends data track very well the data on the number of drivers. Table 2.2 shows that relationship in regression form. Specifications

including city-level fixed effects, month-year fixed effects, and city-specific time trends all give similar answers in terms of the strength of the relationship. Regardless of whether the number of Uber drivers is normalized by the population in the city, a one-percentage-point increase in relative search intensity for Uber is associated with a 1.6% increase in the number of Uber drivers relative to the population. The regression as a whole is highly significant and informs the two-sample two-stage least squares estimation procedure. In my preferred model, I get a prediction of the number of Uber drivers in each city by first estimating (to see the full estimation procedure using GMM and adjustments for robust standard errors, see appendix):

$$\begin{aligned} Drivers_{jt} = \beta_0 + \beta_1 Google\ Activity_{jt} + \beta_2 urate_{jt} + \beta_3 Uber\ launched_{jt} \\ + \beta_4 (months\ since\ Uber\ launched)_{jt} + \mu_j + \alpha_t + \epsilon_{jt}, \end{aligned} \tag{16}$$

where $Drivers_{jt}$ is the number of Uber drivers relative to the population and, again, $Google\ Activity_{jt}$ is the search intensity for Uber by city by month, normalized by total Google searches in a city; $urate_{jt}$ is the unemployment rate in city j in month t; $Uber\ launched_{jt}$ is whether Uber is available in city j in month t as coded from press reports; $(months\ since\ Uber\ launched)_{jt}$ is the number of months since Uber launched in city j in month t; μ_j is a vector of city-level fixed effects; and α_t is a vector of month-year fixed effects. The coefficients from this regression are estimated for the 18 cities from which Uber data are available. The estimate of the share of Uber drivers relative to the population are the fitted values from this regression, $\widehat{Drivers}_{jt}$.

The equation to be estimated next takes the following form:

$$Y_{jt} = \delta_0 + \delta_1 \widehat{Drivers}_{jt} + \delta_2 urate_{jt} + \mu_j + \alpha_t + \epsilon_{jt} \tag{17}$$

where the notation is the same as before and where $\widehat{Drivers}_{jt}$ now represents the fitted values from the previous regression for city t in month j. This methodology is a generaliza-

tion of two-stage least squares (2SLS). However, 2SLS is infeasible in this sample because administrative measures on the number of drivers are available for only 18 cities. I thus run two-sample two-stage least squares. Again I explore city-specific time trends in certain specifications and explore different weights. The results are shown in table 2.3, my preferred regression. Regardless of model chosen, the results suggest that Uber has not had a statistically detectable effect on wages or employment. The inclusion of time trends reverses the negative wage effects in column (1). That nonfinding could be caused by the imprecision of my proxy or the imperfectness of my left-hand-side variables, but all of the analysis suggests there is little evidence that wages and employment have been affected dramatically. The data suggest the possibility that employment in the companies covered by the QCEW might have fallen while wages in the sector rose. This would be consistent with formerly lower paid hourly earners in the industry exiting for higher wages with Uber. In a general equilibrium model, management will respond with higher wages to retain drivers.

One might argue that the uptake of Uber is endogenous and what matters is Uber's presence and length of stay. To address that issue, I estimate a difference-in-difference model of the form:

$$Y_{jt} = \beta_0 + \beta_1(\textit{months since Uber launched})_{jt} + \beta_2 \textit{uratejt} + \mu_j + \alpha_t + \epsilon_{jt} \qquad (18)$$

In this equation, β_0 is a constant, and β_1 is the difference-in-difference estimate of the effect of the presence of Uber in a city on wages and employment.

Table 2.4 presents results of my four measures of employment and wages and different Uber services. Standard errors, as always, are clustered at the city level.[29] The results again show that Uber's length of presence in a city is not associated with significant changes in the wage and employment structure in the taxi and limousine industry. Focusing on more-

[29] I have 70 city clusters – more than the generally accepted level outlined in Cameron and Miller (2015). For the 18-city sample, I use robust standard errors.

reliable wage data from the QCEW, which does not include Uber drivers and which misses many other taxi and limousine drivers who freelance, I see that wages are slightly higher in areas into which Uber has spread. The inclusion of various city-level economic characteristics such as unemployment rate does not affect the results. Overall, the estimates suggest no significant relationship between the presence of Uber in a city and wages or employment in the taxi industry – certainly, no significant decline.

It is known that estimates from this type of difference-in-difference model are not valid under many circumstances if, for example, the assumption of parallel trends is violated (Angrist and Pischke 2008; Bertrand et al. 2002). To help alleviate concerns about differing city-level trends, I run relative time models in table 2.5 for my four dependent variables that show the changes in my key left-hand-side variables in cities, normalized by their time relative to Uber's introduction. In line with the hypothesis that Uber spread to cities based primarily on size, on convenience for Uber, and on idiosyncratic regulation rather than on changing conditions in the local taxi industry, it does not appear as if Uber made its entry decisions based on changes in the taxi industry in the pre-period. The negative estimates in the later period for the employment measures is interesting but could be due to sample selection as the cutoff becomes more stringent. The lack of significance in other regressions and the overall employment pattern in the CPS leads us to think it is not a big factor. Finally, following Berger et al. (2016) I attempt triple differencing, comparing relative wages of taxi and limousine drivers to the city overall wage rate, trucking industries, and other low-skilled industries. Regardless of the comparison group I find no significant effects of Uber's entry on taxi and limousine drivers relative to the comparison group. Tables 2.6 and 2.7 show two of these specifications.

Using both timing-based and popularity-based measures of Uber exposure in the largest US MSAs, I found very little evidence that wages have moved down significantly. I found some evidence for increased employment as reported in the CPS, where my measures include

Uber drivers, with relatively little if any employment changes in the sector covered by the QCEW. I note again that the dataset is biased against finding large employment effects because of the limited coverage of the sector in the two main data sources.

2.7 Conclusion

This paper provides the first systematic evidence – based on the wages and employment of taxi drivers – of the nationwide effect of the rise of Uber. Regardless of which measure of Uber's expansion I use, I find that after Uber's entry into new markets, the wages of taxi and limousine drivers have not declined significantly, increasing slightly in the most reliable dataset. This has happened concurrently with a record number of people reporting employment in driving occupations. These findings are at odds with many claims in the media that Uber and other start-ups are undercutting wages and driving full-time workers out of the industry. I present a wide variety of alternative specifications to probe the robustness of that nonfinding. The results deviate from past experiences in the deregulation of highly controlled industries, including the trucking and airline industries in the 1970s and 1980s. A similar dynamic does not appear to be in play in the taxi and limousine industry in general, wherein drivers seemed to benefit from little rent sharing prior to the entry of TNCs. Notably, previous instances of deregulation took place in heavily unionized industries where workers had a vehicle through which to extract rents from management.

The results, though still early, have implications for current policies toward Uber and other disruptors currently being debated. Whether Uber has hurt drivers is a question that warrants further data collection. City governments have to look more closely at the lease fees that drivers are charged and how that translates to take-home pay for drivers both before and after the advent of Uber. Number of trips completed and number of cabs on the road are not sufficient to diagnose earnings changes. Also, the large number of drivers who have

defected to the Uber platform should be included in any welfare calculations of Uber's effect on taxi and limousine industries. More generally, my findings suggest that large disruptions to industries do not necessarily harm incumbent workers, even if they harm incumbent companies.

References

Abe, M. A., and B. C. Brush. "On the regulation of price and service quality: The taxicab problem." Quarterly Review of Economics and Business, Vol. 16 (1976): 105-111.

Abowd, John M., and Francis Kramarz. "A test of negotiation and incentive compensation models using longitudinal French enterprise data." No. w4044. National Bureau of Economic Research, 1992.

Acemoglu, Daron. "Good jobs versus bad jobs." Journal of Labor Economics 19.1 (2001): 1-21.

Alm, James, and Brian Erard. "Using Public Information to Estimate Self-Employment Earnings of Informal Suppliers." Tulane University Department of Economics Working Papers No. 1517, 2015.

Angrist, Joshua D., and Jörn-Steffen Pischke. "Mostly harmless econometrics: An empiricist's companion." Princeton University Press, 2008.

Ashenfelter, Orley, Kirk Doran, and Bruce Schaller. "A Shred of Credible Evidence on the Long-Run Elasticity of Labour Supply." Economica 77.308 (2010): 637-650.

Baanders, Ambrosius, and Marcel Canoy. "Ten Years of Taxi Deregulation in the Netherlands—the Case for Re-regulation and Decentralisation." European Transport Conference, 2010.

Barrett, Sean D. "The sustained impacts of taxi deregulation." Economic Affairs 30.1 (2010): 61-65.

Berger, Thor, Chinchih Chen, and Carl Benedikt Frey. "Drivers of disruption? Estimating the Uber effect." Working paper, 2017.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. "How much should we trust differences-in-differences estimates?" No. w8841. National Bureau of Economic Research, 2002.

Blanchflower, David G., Andrew J. Oswald, and Peter Sanfey. "Wages, profits and rent-sharing." No. w4222. National Bureau of Economic Research, 1992.

Buchholz, Nicholas. "Spatial equilibrium, search frictions and efficient regulation in the taxi industry." Working Paper, University of Texas at Austin, 2016.

Bureau of Labor Statistics. US Department of Labor, Occupational Employment Statistics, www.bls.gov/oes/.

Cairns, Robert D., and Catherine Liston-Heyes. "Competition and regulation in the taxi industry." Journal of Public Economics 59.1 (1996): 1-15.

Camerer, C., et al. "Labor supply of New York taxi drivers: One day at a time." Quarterly Journal of Economics 112 (1997): 407-441.

Cameron, A. Colin, and Douglas L. Miller. "A practitioner's guide to cluster-robust inference." Journal of Human Resources 50.2 (2015): 317-372.

Card, David. "Deregulation and Labor Earnings in the Airline Industry." In James Peoples, editor, Regulatory Reform and Labor Markets. Norwell, MA: Kluwer Academic Publishers, 1998.

Choi, Hyunyoung, and Hal Varian. "Predicting the present with Google Trends." Economic Record 88.s1 (2012): 2-9.

Cohen, P., Hahn, R., Hall, J., Levitt, S., & Metcalfe, R. (2016). "Using big data to estimate consumer surplus: The case of Uber" (No. w22627). National Bureau of Economic Research. 2016

Cramer, Judd, and Alan B. Krueger. "Disruptive change in the taxi business: The case of Uber." The American Economic Review 106.5 (2016): 177-182.

Coffman, R. B. "The economic reasons for price and entry regulation of taxicabs: A comment." Journal of Transport Economics and Policy 11(3) (1977): 288-297.

Dempsey, Paul Stephen. "The rise and fall of the Interstate Commerce Commission: The tortuous path from regulation to deregulation of America's Infrastructure." Marquette Law Review 95 (2012): 1151-1679.

De Vany, A. S. "Capacity utilization under alternative regulatory constraints: An analysis of taxi markets." Journal of Political Economy 83(1) (1975): 83-94.

Douglas, G. W. "Price regulation and optimal service standards: The taxicab industry." Journal of Transport Economics and Policy 20 (1972): 116-127.

Edelman, Benjamin G., and Damien Geradin. "Efficiencies and regulatory shortcuts: How should we regulate companies like Airbnb and Uber?" Harvard Business School NOM Unit Working Paper 16-026 (2015).

Farber, Henry S. "Reference-dependent preferences and labor supply: The case of New York City taxi drivers." American Economic Review 98.3 (2008): 1069-1082.

Farber, Henry S. "Why you can't find a taxi in the rain and other labor supply lessons from cab drivers." No. w20604. Quarterly Journal of Economics, 2015.

Frankena, Mark W., and Paul A. Pautler. "Taxicab regulation: An economic analysis." Research in Law and Economics 9 (1986): 129-165.

Frechette, Guillaume R., Alessandro Lizzeri, and Tobias Salz. "Frictions in a Competitive, Regulated Market: Evidence from Taxis." October 19, 2015.

Goodbody Economic Consultants. "Economic Review of the Small Public Service Vehicle Industry." Dublin: Commission for Taxi Regulation, 2009.

Greenwood, Brad N., and Sunil Wattal. "Show Me the Way to Go Home: An Empirical Investigation of Ride Sharing and Alcohol Related Motor Vehicle Homicide." Available at SSRN 2557612 (2015).

Hall, Jonathan V., and Alan B. Krueger. "An Analysis of the Labor Market for Uber's Driver-Partners in the United States." Princeton University Working Papers (2015).

Hall, Jonathan V., Cory Kendrick, and Chris Nosko. The Effects of Uber's Surge Pricing: A Case Study. Working paper (2015).

Harper, Donald V. "The Federal Motor Carrier Act of 1980: Review and analysis." Transportation Journal (1980): 5-33.

Hirsch, Barry T. "Trucking deregulation and labor earnings: Is the union premium a compensating differential?" Journal of Labor Economics (1993): 279-301.

Hurst, Erik, Geng Li, and Benjamin Pugsley. "Are household surveys like tax forms? Evidence from income underreporting of the self-employed." Review of Economics and Statistics 96.1 (2014): 19-33.

Kleiner, Morris M. "Occupational licensing." Journal of Economic Perspectives (2000): 189-202.

Kuttner, Robert. "The Task Rabbit Economy." Web blog post. American Prospect, October 10, 2013; Web, November 10, 2015.

Lanzona, Leonardo, and Robert E. Evenson. "The effects of transaction costs on labor market participation and earnings: Evidence from rural Philippine markets." Yale Economic Growth Center Disc. Paper 790 (1997).

Manski, Charles F., and J. David Wright. "Nature of equilibrium in the market for taxi services." No. Intrm Rpt. (1967).

New York City Taxi and Limousine Commission. 2014 Annual Report. Peoples, James. "Deregulation and the labor market." Journal of Economic Perspectives (1998): 111-130.

Reich, Robert. "The Share-the-Scraps Economy." Web blog post. Roberreich.org, February 2, 2015; Web, February 2, 2015.

Rose, Nancy L. "Labor rent sharing and regulation: Evidence from the trucking industry." Journal of Political Economy (1987): 1146-1178.

Schneiderman, Todd. "Analyzing 1.1 Billion NYC Taxi and Uber Trips, with a Vengeance." Personal Web site. November, 18, 2015.

Simpson, Hilery. "Tips and Excluded Workers: The New Orleans Test." Compensation and Working Conditions. Bureau of Labor Statistics. Summer 1997.

Stephens-Davidowitz, Seth. "The cost of racial animus on a black candidate: Evidence using Google search data." Journal of Public Economics 118 (2014): 26-40.

Svenjar, Jan. "On the Empirical Testing of the Nash-Zeuthen Bargaining Solution," Industrial and Labor Relations Review, 33, (1980): 536-542.

Teal, Roger F., and Mary Berglund. "The impacts of taxicab deregulation in the USA." Journal of transport economics and policy (1987): 37-56.

Tullock, G. "The transitional gains trap." Bell Journal of Economics 6(2) (1975): 67-78.

Weise, Karen. "How Uber takes over a city." Bloomberg Businessweek, June 23, 2015.

Williams, D. J. "The economic reasons for price and entry regulation of taxicabs: A comment." Journal of Transport Economics and Policy 14(1) (1980): 105-112.

Tables

Table 2.1: Uber effect on wages and employment in Uber's 18 largest markets (2008 to 2014)

	(1) Log weekly wage QCEW	(2) Log hourly wage CPS	(3) Log employment QCEW	(4) Log employment CPS
Log(Active Uber Drivers per Person)	0.001	-0.017	-0.008	-0.008
	(0.004)	(0.032)	(0.008)	(0.029)
Unemployment Rate	0.015	0.251	-0.522	-2.086
	(0.321)	(3.223)	(0.560)	(2.129)
Constant	6.034***	2.328***	6.642***	8.947***
	(0.071)	(0.573)	(0.143)	(0.447)
Observations	467	679	1,373	1,343
R-squared	0.925	0.173	0.993	0.804
City FE	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES

Notes: All wages are measured in real 2016 dollars and represent city-level averages. Average weekly wages in the QCEW are available quarterly and all other outcomes are available monthly. Observations are weighted by the size of the civilian labor force. I add one to the number of drivers before dividing by population over 15 and taking the logarithm. Robust standard errors in parentheses (*** p<0.01, ** p<0.05, * p<0.1).

Table 2.2: Relationship between Uber drivers and Uber-related Google search activity

	Log(Active Uber Drivers per Person)
Log(Uber Google Search Intensity)	1.616***
	(0.348)
Constant	-15.056***
	(0.238)
Observations	1,116
R-squared	0.886

Notes: Data on the number of Uber drivers are available for Uber's 18 largest markets. I add one to the number of drivers before dividing by population over 15 and taking the logarithm. Robust standard errors in parentheses (*** p<0.01, ** p<0.05, * p<0.1).

Table 2.3: Two-sample Two Stage Least Squares estimates of popularity of Uber on taxi and limousine wages and employment

	(1) Log weekly wage QCEW	(2) Log weekly wage QCEW	(3) Log hourly wage CPS	(4) Log hourly wage CPS	(5) Log employment QCEW	(6) Log employment QCEW	(7) Log employment CPS	(8) Log employment CPS
Predicted log drivers to population from first stage	0.014	0.021**	0.005	-0.002	-0.004	-0.041*	-0.046	-0.020
	(0.014)	(0.009)	(0.055)	(0.119)	(0.021)	(0.024)	(0.052)	(0.085)
Indicator for Uber launch	-0.058	-0.063**	-0.075	-0.067	-0.003	0.084	0.131	0.050
	(0.044)	(0.029)	(0.187)	(0.313)	(0.060)	(0.053)	(0.193)	(0.221)
Months since Uber launch	-0.000	-0.001	-0.007	-0.008	0.001	-0.002	0.012	0.005
	(0.002)	(0.003)	(0.009)	(0.014)	(0.004)	(0.003)	(0.010)	(0.010)
Unemployment rate	-0.247	-0.181	3.672	4.661	0.031	0.345	-0.338	-1.028
	(0.191)	(0.184)	(2.586)	(2.848)	(0.422)	(0.260)	(1.422)	(1.381)
Constant	6.123***	5.075***	2.143**	-5.914**	4.841***	9.873***	6.905***	10.479***
	(0.226)	(0.212)	(1.005)	(2.620)	(0.343)	(0.308)	(0.888)	(1.218)
Observations	1,246	1,246	1,214	1,214	3,738	3,738	2,926	2,926
R-squared	0.876	0.912	0.194	0.226	0.989	0.995	0.814	0.823
City FE	YES	YES	YES	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES	YES	YES	YES
City-time trend	NO	YES	NO	YES	NO	YES	NO	YES

Notes: All wages are measured in real 2016 dollars and represent city-level averages. Average weekly wages in the QCEW are available quarterly and all other outcomes are available monthly. Observations are weighted by the size of the civilian labor force. Predicted log of drivers is obtained from a first stage regression of log drivers per population on search intensity and all other covariates. Robust standard errors in parentheses (*** p<0.01, ** p<0.05, * p<0.1).

Table 2.4: Difference-in-Difference estimates of the effect of the length of Uber's presence on employment and earnings of limousine and taxi drivers

	(1) Log weekly wage QCEW	(2) Log weekly wage QCEW	(3) Log hourly wage CPS	(4) Log hourly wage CPS	(5) Log employment QCEW	(6) Log employment QCEW	(7) Log employment CPS	(8) Log employment CPS
Months since Uber launched	0.002	0.002	-0.002	-0.014*	0.000	0.001	0.002	0.005
	(0.002)	(0.001)	(0.004)	(0.007)	(0.004)	(0.002)	(0.005)	(0.007)
Unemployment Rate	-0.174	-0.223	1.215	3.086	0.162	0.211	-2.025*	-1.217
	(0.181)	(0.168)	(1.807)	(1.961)	(0.430)	(0.217)	(1.208)	(1.099)
Constant	5.976***	5.211***	2.655***	0.332	5.061***	8.462***	7.452***	9.774***
	(0.017)	(0.081)	(0.288)	(0.458)	(0.043)	(0.150)	(0.096)	(0.344)
Observations	1,593	1,593	2,008	2,008	4,779	4,779	4,908	4,908
R-squared	0.866	0.903	0.180	0.200	0.987	0.994	0.818	0.827
City FE	YES	YES	YES	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES	YES	YES	YES
City-time trend	NO	YES	NO	YES	NO	YES	NO	YES

Notes: All wages are measured in real 2006 dollars and represent city-level averages. Average weekly wages in the QCEW are available quarterly and all other outcomes are available monthly. Observations are weighted by the size of the civilian labor force. Robust standard errors in parentheses (*** p<0.01, ** p<0.05, * p<0.1).

Table 2.5: Relative time models for Uber

	(1) Log weekly wage QCEW	(2) Log hourly wage CPS	(3) Log employment QCEW	(4) Log employment CPS
Eight Quarters Before	0.00199	-0.272	-0.009	-0.046
	(0.00907)	(0.234)	(0.019)	(0.114)
Seven Quarters Before	-0.0148	-0.220	-0.001	0.168
	(0.0130)	(0.154)	(0.023)	(0.113)
Six Quarters Before	-0.00704	0.016	-0.008	-0.049
	(0.0117)	(0.090)	(0.026)	(0.149)
Five Quarters Before	-0.00892	0.031	0.005	0.079
	(0.00995)	(0.081)	(0.028)	(0.098)
Four Quarters Before	0.0103	0.217	-0.005	0.009
	(0.0117)	(0.191)	(0.027)	(0.085)
Three Quarters Before	-0.00168	-0.068	0.000	-0.099
	(0.0114)	(0.087)	(0.027)	(0.104)
Two Quarters Before	-0.00552	0.001	-0.005	-0.092
	(0.0122)	(0.135)	(0.026)	(0.119)
One Quarter Before	-0.0215*	0.132	-0.001	-0.111
	(0.0108)	(0.093)	(0.022)	(0.085)
One Quarter After	-0.0113	0.068	-0.006	-0.035
	(0.0123)	(0.120)	(0.018)	(0.073)
Two Quarters After	-0.0274*	0.070	-0.004	-0.121
	(0.0139)	(0.042)	(0.021)	(0.107)
Three Quarters After	-0.0279*	-0.372**	-0.023	0.051
	(0.0145)	(0.169)	(0.024)	(0.085)
Four Quarters After	-0.0286**	-0.038	-0.027	-0.037
	(0.0115)	(0.096)	(0.026)	(0.077)
Five Quarters After	-0.0227*	0.050	-0.032	-0.068
	(0.0123)	(0.113)	(0.026)	(0.079)
Six Quarters After	-0.0166	0.062	-0.024	-0.235*
	(0.0131)	(0.094)	(0.026)	(0.122)
Seven Quarters After	-0.0363**	0.046	-0.033	-0.177**
	(0.0175)	(0.106)	(0.029)	(0.079)
Eight Quarters After	-0.0234	-0.058	-0.025	-0.351***
	(0.0155)	(0.144)	(0.026)	(0.126)
Constant	6.117***	3.750***	5.649***	7.529***
	(0.0155)	(0.368)	(0.040)	(0.114)
Observations	1,784	3,229	5,352	9,138
R-squared	0.883	0.160	0.987	0.852

Notes: All wages are measured in real 2016 dollars and represent city-level averages. Average weekly wages in the QCEW are available quarterly and all other outcomes are available monthly. Observations are weighted by the size of the civilian labor force. Omitted time period is quarter of Uber launch. Robust standard errors in parentheses (*** p<0.01, ** p<0.05, * p<0.1).

Table 2.6: The effect of Uber on taxi and limousine wages and employment using truck drivers as a control group

	(1) Relative log weekly wage QCEW	(2) Relative log weekly wage QCEW	(3) Relative log weekly wage CPS	(4) Relative log weekly wage CPS	(5) Relative log employment QCEW	(6) Relative log employment QCEW	(7) Relative log employment CPS	(8) Relative log employment CPS
Log(Uber Google Search Intensity)	0.008	-0.013	-0.005	0.064	0.007	0.007	-0.006	0.019
	(0.010)	(0.011)	(0.084)	(0.104)	(0.023)	(0.016)	(0.057)	(0.059)
Unemp. Rate	0.024	0.005	5.072	4.053	0.316	0.213	0.258	0.362
	(0.185)	(0.198)	(3.056)	(3.174)	(0.454)	(0.265)	(1.879)	(1.461)
Constant	-0.733***	-1.086***	-2.832***	-72.833***	-2.135***	-7.259***	-0.571***	-0.434
	(0.023)	(0.111)	(0.406)	(11.786)	(0.041)	(0.273)	(0.211)	(0.951)
Observations	1,175	1,175	1,146	1,146	3,525	3,525	3,163	3,163
R-squared	0.880	0.915	0.167	0.213	0.979	0.990	0.531	0.566
City FE	YES	YES	YES	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES	YES	YES	YES
City-time trend	NO	YES	NO	YES	NO	YES	NO	YES

Notes: The relative log wage in the log of taxi drivers' wages divided by truck drivers' wages. Similarly, relative log employment is the log of taxi employment divided by trucking employment. All wages are measured in real 2016 dollars and represent city-level averages. Average weekly wages in the QCEW are available quarterly and all other outcomes are available monthly. Observations are weighted by the size of the civilian labor force. Robust standard errors in parentheses (*** p<0.01, ** p<0.05, * p<0.1).

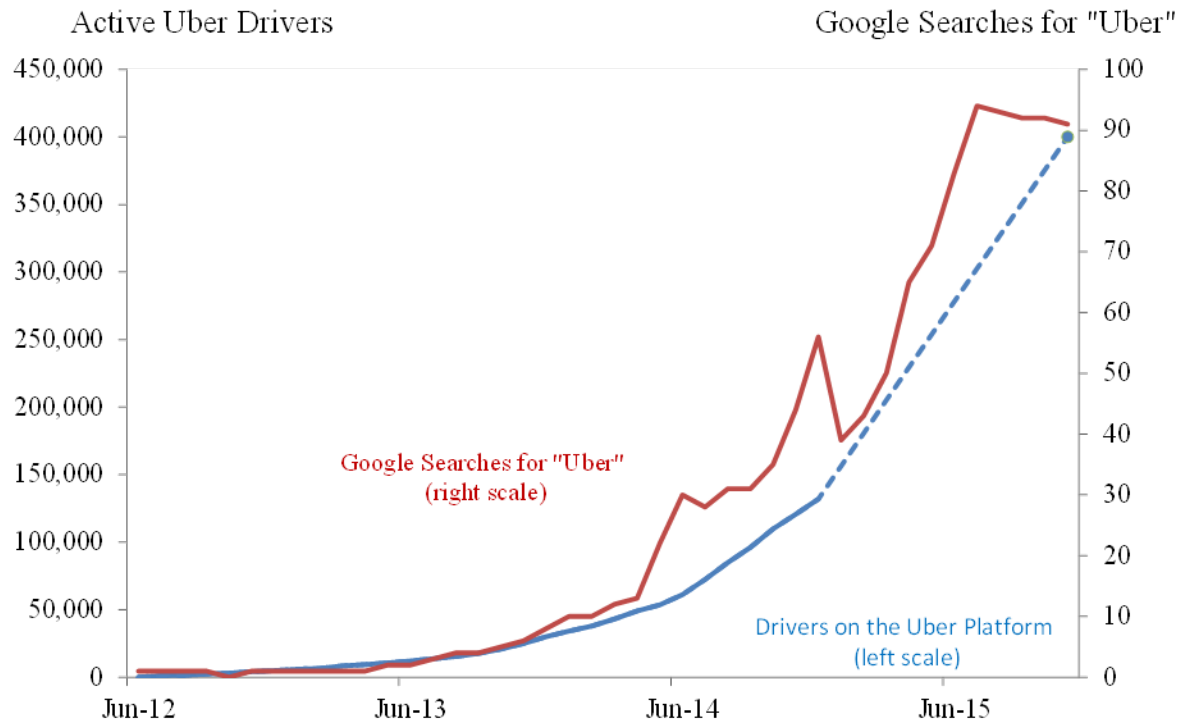
Table 2.7: The effect of Uber on taxi and limousine wages and employment using all other employment as a control group

	(1) Relative log weekly wage QCEW	(2) Relative log weekly wage QCEW	-3 Relative log weekly wage CPS	(4) Relative log weekly wage CPS	(5) Relative log employment QCEW	(6) Relative log employment QCEW	(7) Relative log employment CPS	(8) Relative log employment CPS
Log(Uber Google Search Intensity)	0.005	-0.012	0.036	0.115	-0.003	0.004	0.005	0.013
	(0.012)	(0.014)	(0.075)	(0.087)	(0.018)	(0.014)	(0.046)	(0.049)
Unemp. Rate	-0.278	-0.319*	5.141*	3.967	0.172	0.152	-0.017	0.155
	(0.185)	(0.191)	(2.834)	(2.904)	(0.405)	(0.250)	(1.472)	(1.181)
Constant	-0.811***	-0.415	-2.389***	-46.121***	-6.979***	-11.619***	-4.610***	-6.930***
	(0.027)	(0.381)	(0.372)	(4.804)	(0.036)	(0.330)	(0.175)	(0.986)
Observations	1,363	1,363	1,283	1,283	4,089	4,089	3,190	3,190
R-squared	0.802	0.849	0.168	0.212	0.956	0.977	0.512	0.540
City FE	YES	YES	YES	YES	YES	YES	YES	YES
Month-Year FE	YES	YES	YES	YES	YES	YES	YES	YES
City-time trend	NO	YES	NO	YES	NO	YES	NO	YES

Notes: The relative log wage in the log of taxi drivers' wages divided by wages of workers in all other occupations. Similarly, relative log employment is the log of taxi employment divided by all other employment. All wages are measured in real 2016 dollars and represent city-level averages. Average weekly wages in the QCEW are available quarterly and all other outcomes are available monthly. Observations are weighted by the size of the civilian labor force. Robust standard errors in parentheses (*** p<0.01, ** p<0.05, * p<0.1).

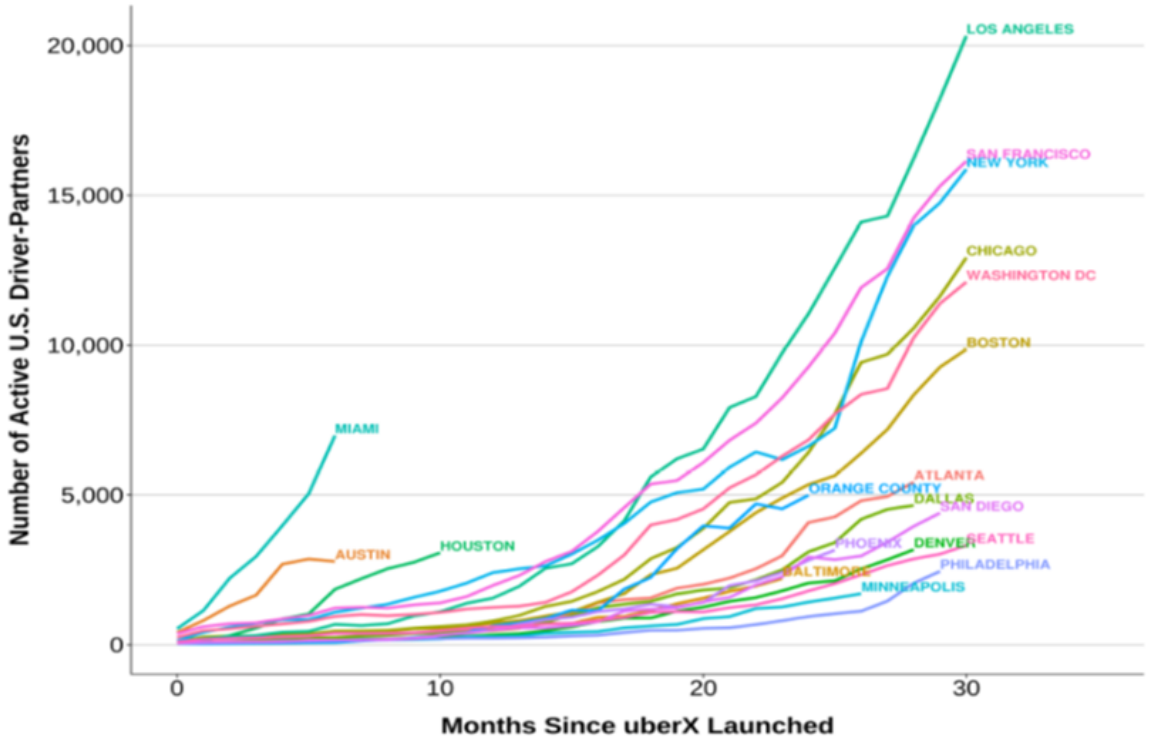
Figures

Figure 2.1: Drivers on the Uber System in the US and Google Searches for "Uber", Jun. 2012 to Oct. 2015



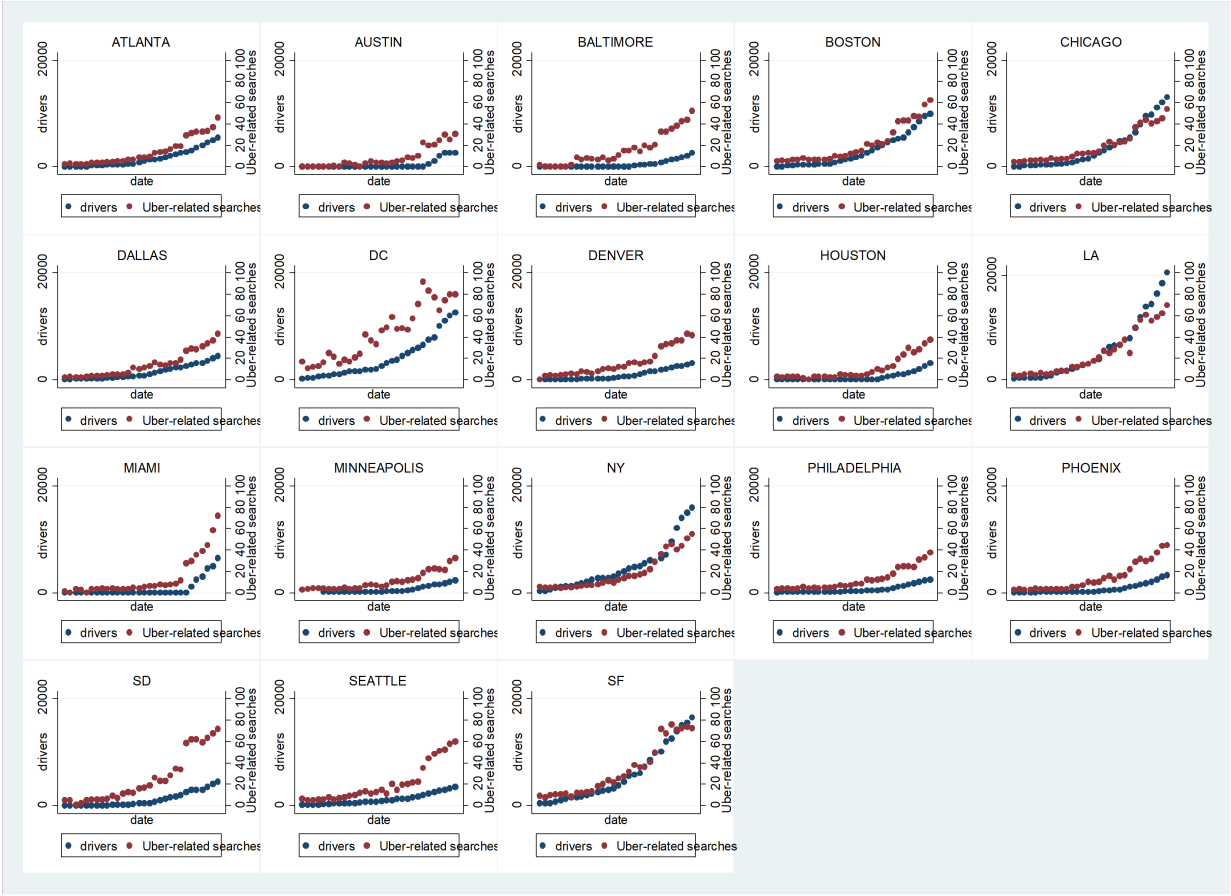
Source: Google Trends, Hall and Krueger (2015), https://newsroom.uber.com/2015/11/1776/.
Notes: Dashed-line represents linear interpolation between number of active drivers released in Hall and Krueger (2015) and David Plouffe's statement that there were 400,000 active drivers on November 3, 2015; Google Trends data normalized to zero in October 2012.

Figure 2.2: Active Uber Drivers by City, Jun. 2012- Dec. 2014 (Hall and Krueger, 2015)



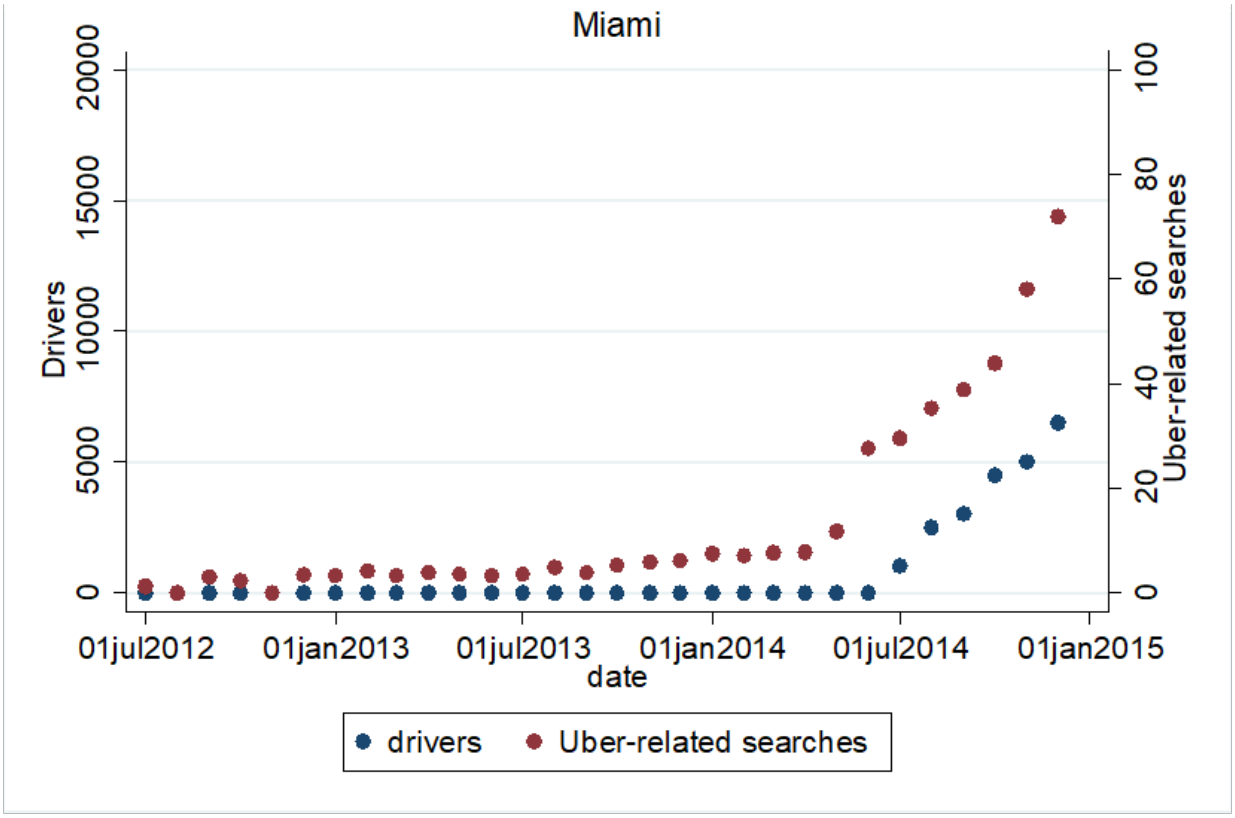
Notes: From Hall and Krueger (2015): "Figure reports the number of U.S. UberBLACK and uberX driver-partners making at least one trip in the specified month, indexed to the number of months since Uber began in the city or June 2012, whichever came first."
Source: Hall and Krueger (2015)

Figure 2.3: Uber Drivers vs. Uber Search Intensity



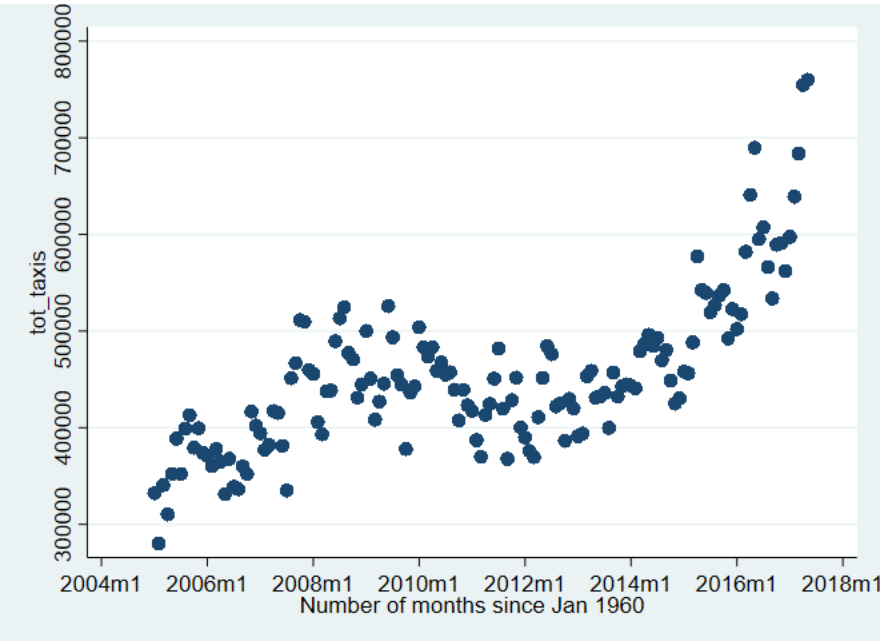
Source: Google Trends, Krueger and Hall (2015), author's calculations.

Figure 2.4: Uber Drivers vs. Uber-related Google Searches



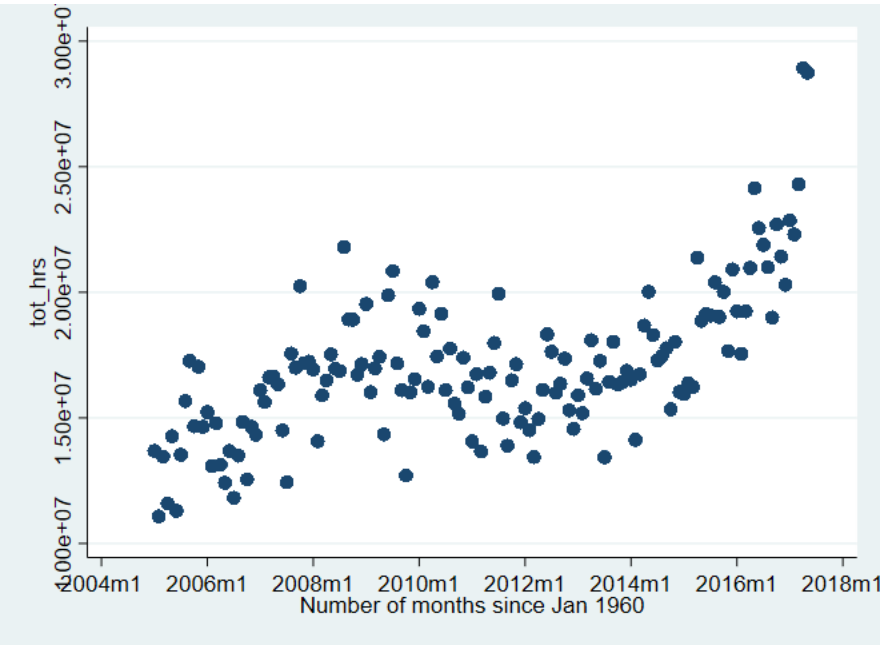
Source: Google Trends, Krueger and Hall (2015), author's calculations.

Figure 2.5: Total Taxi Employment, Jan. 2008-Jun. 2017 (Current Population Survey)



Source: Current Population Survey, Author's Calculations

Figure 2.6: Total Taxi and Limousine Hours Worked, Jan. 2008-Jun. 2017 (Current Population Survey)



Source: Current Population Survey, Author's Calculations.

Chapter 3

Long-term Unemployment and the Loop in the Beveridge Curve: The Importance of Cyclical Labor Force Withdrawal by Duration

3.1 Introduction

One of the most noteworthy and distressing features of the Great Recession and its aftermath in the United States was the unprecedented increase in long-term unemployment (Farber, 2015). By April 2014, the rate of long-term unemployment was still more than double its pre-recession average; in contrast, the number of short-term unemployed workers relative to the labor force had nearly returned to normal levels (Krueger et al., 2014). At this time, many of the seminal relationships in the macroeconomy between the unemployment rate and other variables appeared to have partially or completely broken down (e.g., Okun's Law, the Phillips curve, and the Beveridge curve).[1] In particular, in the beginning of 2014, the Beveridge curve, which traces the relationship between the unemployment rate and the job vacancy rate appeared to have shifted out relative to the pre-recession period, with the unemployment rate elevated around 2.0 percentage points (p.p.) above its previous value

[1] See for example, IMF (2010) on Okun's law, Gordon (2013) and Stock (2011) on respecifying the Phillips Curve, and Krueger et al. (2014) on the Beveridge curve.

when the job opening rate was at similar levels (see Figure 3.1(a)).[2] When measures of short-term unemployment were used, however, the Beveridge curve appeared stable, experiencing movements along the curve as opposed to outward shifts (see Figure 3.1(b)). It was unclear, therefore, if the U.S. economy was in the process of beginning a counterclockwise loop around a stable Beveridge curve as Blanchard and Diamond (1989) outlined in their canonical paper, or if the overall relationship had simply changed. Various academics suggested that skills mismatch, secular stagnation, or declines in recruiting intensity by employers could all have fundamentally altered the Beveridge curve relationship.[3] This presented a puzzle for policymakers about how much slack there was in the economy and the relative efficacy of fiscal and monetary policy on the unemployment rate, which many argued was at or near its natural rate by mid-2014.[4]

This paper documents that it was the cyclical behavior of the record number of long-term unemployed during and after the recession that is crucial to explaining this puzzle. The labor market decisions of the long-term unemployed explains both the initial movements away from the Beveridge curve and the eventual completion of a loop (as shown in Figure 3.2). To estimate the labor market effects of the long-term unemployed throughout the business cycle precisely, I derive and expand upon the model originally calibrated in Krueger, Cramer, and Cho (KCC) (2014), whose forecast has performed well throughout the recovery. In particular, I examine the channels through which duration dependence for the long-term unemployed was a factor. The goal of the simulation model is to quantify the effect of the long-term unemployed on the economy since the beginning of the Great Recession. Incorporating the differing labor market behavior throughout a jobless span improved the

[2]Barnichon et al. (2011) show that the Beveridge curve appeared to have shifted out by 2.3 p.p. in January 2011.

[3]See for example Daly et al. (2011) on possible shifts in the Beveridge curve, Davis et al. (2012) on recruiting intensity or Summers (2014) for background on the secular stagnation hypothesis.

[4]The unemployment rate stood at 6.4 percent in April 2014. The Federal Reserve in that month predicted that the unemployment rate would decline gradually to 5.1 to 5.8 percent in 2016 and settle between 5.2 and 6.0 percent in the "longer run." March 2014 Federal Reserve forecast, accessible at: https://www.federalreserve.gov/monetarypolicy/files/fomcprojtabl20140319.pdf.

accuracy of the model's predictions of the unemployment rate during the recession by 60 percent and resulted in unemployment rate forecasts which were 15 percent more accurate than the Federal Reserve and the Survey of Professional Forecasters from 2014 to 2016.

With the recovery having lasted nine years by the middle of 2018, enough time has elapsed to take stock of what has been learned about the record number of long-term jobless during and following the recession. Relative pessimism about the economic futures of the long-term jobless has proven accurate. The job-finding rate for the long-term unemployed and nonparticipants has remained at depressed levels, due to their lower search- and match-efficiency even during expansions. On the other hand, the labor force withdrawal rate for the long-term unemployed has increased rapidly, in line with past patterns during recoveries; although this was not widely acknowledged previously. The labor force reentry rates of those who had exited the labor force during and after the recession has remained depressed. This stands in contrast to the predictions made by many commentators, notably Goldman Sachs and Federal Reserve Chair Janet Yellen, who suggested that the improving economy could lead to an influx of non-attached workers which could push the labor force participation rate higher in spite of demographic headwinds.[5] As of current writing, there has still been no increase in labor force reentry for the previously long-term unemployed. In sum, the situation for the long-term jobless has remained strikingly similar to the depressing evidence KCC compiled in 2014. They showed then using the Current Population Survey (CPS), the Survey of Income and Program Participation (SIPP), and other sources, that marginal workers apply for fewer jobs over their nonemployment spells, get called back for interviews for the jobs they do apply for at a lower rate, suffer from depression and other mental health issues at higher rates as their unemployment spells lengthen, and have acyclical labor

[5]See for example CEA report "The Labor Force Participation Rate Since 2007: Causes and Policy Implications," which attributed half of the decline in LFPR to aging. Demographic changes could have also included people who were going back to school or receiving other retraining. The labor force flows of prime-aged men are strikingly similar.

force reentry rates.[6] As the recovery has continued, these workers have kept struggling in the same ways, and, as this paper documents, their actions have had large effects on the macroeconomy.

The rest of the paper is structured as follows. Section 2 gives background on the recent labor market, especially the performance of the long-term jobless, drawing evidence mainly from the CPS. Section 3 derives and expands the model first introduced in KCC and explores further simulation exercises on the Beveridge curve. Section 4 summarizes and concludes.

3.2 Update on Those on the Margins of the Labor Force

3.2.1 Labor Force Performance of the Unemployed

For background, this section provides a detailed description of the economic performance of the long-term unemployed and labor force nonparticipants during the recovery by examining the labor force transitions of those in the CPS over two consecutive months and longer periods. Figure 3.3 shows that for most of the six decades prior to the Great Recession, the share of the unemployed in the United States who were out of work for more than half a year oscillated between 10 and 20 percent during both recoveries and recessions, although there had been upward drift in this series and median duration more generally over this period.[7] The share of the unemployed who were long-term job seekers averaged over 40 percent in 2010, 2011 and 2012—reaching as high as 45 percent—and remained at 40 percent in April 2014. The long-term unemployed, therefore, had a more significant effect on unemployment

[6]The original research on search activity callbacks comes from Krueger and Mueller (2011) and Wanberg et al. (2012), callback rates from Kroft, Lange, and Notowidigdo (2013) and Ghayad (2013), Krueger (2017) discusses life satisfaction and opioid abuse in greater depth.

[7]See Abraham and Shimer (2002) for a careful analysis of the three main reasons why the duration of unemployment in the U.S. generally rose relative to the unemployment rate in the 1980s and 1990s: the CPS redesign of 1994; the aging of the baby-boomers; and, a decline in the unemployment rate for women. These changes led to higher shares classified as long-term unemployed in a given month.

dynamics in this recovery than they have in the past and continued to have outsized influence on the labor market into 2018, when the long-term share still stood at 20 percent.

KCC noted that the characteristics of the long-term unemployed in early 2014 would have predicted more beneficial outcomes than previous cohorts of long-term unemployed.[8] In spite of this, Figure 3.4 shows using flow data from the CPS that the job-finding rate for the long-term unemployed has barely returned to its pre-recession average, a disappointing performance when considering that the unemployment rate stands at 4.0 percent and the expansion has continued since July 2009.[9] Figure 3.5 further disaggregates the duration groups and shows that the subset of the long-term unemployed who have been out of work for over a year are still faring much worse than they had been before the recession and look very similar to how they performed during the depths of the recession.[10] Even during the economic boom of the late 1990s, the long-term unemployed fared significantly worse than other groups, and that has continued during the recovery. Figure 3.6 shows, using the same data source, that while the long-term unemployed have not been finding employment, the probability of reporting that they are out of the labor force has continued to increase. Figure 3.7 again further disaggregates and shows that those who have been out of work for over a year are leaving the labor force at record levels. The pattern looks like what KCC predicted when they said: "after labor force exit rates collapse in recessions to about the same level

[8]See for example, the discussion around figures 3.3-3.9, that showed that the long-term unemployed were projected to have higher job-finding rates and wages, and lower labor force withdrawal rates based on regressions run in the pre-recession period. This pattern is consistent with Andreas Mueller's (2012) finding that in recessions the pool of the unemployed tends to shift toward those with higher earnings in their previous jobs, because such workers are more likely to be displaced in recessions.

[9]Many researchers have documented that CPS data can severely misstate gross labor market flows because of classification errors, which means one must use caution in interpreting the data. For instance, studies based on re-interview data from the CPS in which respondents are re-interviewed about one week after their initial interview have found that a person's reported labor force status can be wildly inconsistent. Nevertheless, we believe there is signal in this series and it has advantages coming from the survey that is used to calculate the unemployment rate. See Abowd and Zellner (1985) and Poterba and Summers (1986) for further discussion.

[10]In this section, we focus on overall rates. The figures remain nearly identical if one focuses solely on prime-aged workers or prime-aged men alone to abstract from increased schooling and the aging of the population.

for all duration groups, they tend to move back in recoveries toward their historical norms with a higher exit rate for the long-term unemployed."[11] Taken in concert with Figure 3.5, in 2018, somebody who has been unemployed for over a year is over three times more likely to have left the labor force than to have found a job in their next survey month. These labor force transitions for the long-term unemployed led to the share of unemployed who were long-term unemployed to fall back within its pre-recession range at 20 percent in April 2018, though it is still well above previous lows (see Figure 3.3). The monthly flows for the short-term unemployed, in contrast, have all recovered to look roughly in line with past levels. As shown in the simulations below based on a job-matching model calibrated to data from 2002 to 2007, the movement of the labor force withdrawal rates of the long-term unemployed toward their historical averages has important implications for the unemployment rate and, relatedly, for the Beveridge curve.

Although the matching model calibrated and simulated in Section 3 is based only on estimates of monthly transitions from the CPS, to verify that the long-term unemployed have remained on the margins of the labor market since, I also look at transition rates over longer periods of time. In the CPS's rotation group design, individuals are interviewed for 4 consecutive months, then dropped out of the survey for 8 months, and then interviewed again for 4 more months. This design makes it possible to examine transitions over a 15-month interval. A monthly job finding rate of 10 to 15 percent for the long-term unemployed would exaggerate their connection to employment if random classification errors in labor force status inflate their transition rates or if the long-term unemployed tend to work in transitory jobs if they do become reemployed, as Ann Stevens (1997) finds. One way to assess the importance of these issues, and obtain an alternative indicator of the extent to which the long-term unemployed are connected to the labor market, is to examine transition

[11] This phenomenon probably reflects, in part, the extension and subsequent exhaustion of unemployment insurance benefits, which require workers to search for a job and have been found to induce unemployed workers to stay in the labor force, thus elevating the measured unemployment rate and partially elevating exit rates afterwards (see Farber and Valletta, 2015).

rates over a longer period of time. I can then compare how well those who were long-term unemployed in 2014 and later wound up doing after 15 months. Indeed, the actual long-term transition rates are considerably lower than those implied by monthly data because many workers transition across categories many times.[12] The results again are discouraging for the long-term unemployed, with no recovery seen since the writing of KCC. Alarmingly, since the beginning of the Great Recession, the probability of being employed 15 months later for those who were long-term unemployed in their first month in the CPS has yet to recover to its pre-recession average and remains at just 31 percent in 2016.

3.2.2 Labor Force Performance of Nonparticipants

In her August 2014 speech (Yellen, 2014), Janet Yellen said that "we could reasonably expect further increases in labor demand to pull a sizable share of discouraged workers back into the workforce." Figure 3.8 shows that the transition rates for those out of the labor force have remained depressed, with no bounce back in the recovery. A larger share of nonparticipants move into employment than into unemployment, consistent with Barnichon and Figura (2013), but that flow remains below its pre-recession average. The longer run transitions of nonparticipants are also bleak. People who reported being out the labor force initially in the CPS, had an 86 percent probability of reporting being out of the labor force 15-months later for every year from 2009 to 2015, when rounded to the nearest whole unit. In 2016, it bumped up to 87 percent. There has been virtually no movement in this series throughout the recovery. Of those who transition into the labor force after 15 months, 10 percent find employment and 4 percent move to unemployment on average since 2012. Overall, the performance of the long-term unemployed and nonparticipants has been as bad as one could have feared, with the ongoing recovery failing to undo the losses from the recession.

[12] If there were independence and a constant 0.10 probability of finding a job in any given month and jobs were permanent, the proportion of unemployed workers who would have gained employment within 15 months would be $1-(1-0.10)^{*}15=0.79$.

3.3 The Beveridge Curve

After the Great Recession the relationship between the vacancy rate and the unemployment rate, known as the Beveridge curve, appeared to shift outward, with more vacancies than predicted given the high unemployment rate (see Figure 3.1(a)).[13] Some argued that either the Job Creation Curve or the Beveridge curve had shifted and the natural rate of unemployment had increased.[14] Alternate theories posited that the meaning of a vacancy had changed or that skill mismatch made some of the long-term unemployed unhireable.[15] If any of these explanations were accurate, they would not necessarily predict a return to the previous Beveridge curve. Another possibility was that nothing had fundamentally changed, and the economy would return to its original position. This could happen after completing a counterclockwise loop generated in Blanchard and Diamond (1989) from random aggregate shocks that followed a sine wave. Similarly, if employers rank job applicants based on their duration of unemployment when they receive multiple applicants, long-term unemployment could persist longer than it would otherwise (Blanchard and Diamond (1994)). This lack of job-finding coupled with a decline in labor force exits, particularly among the long-term unemployed as documented in section 2 could then return the economy to the original Beveridge curve. The goal in this section is to explore how these hypotheses could be tested with a calibrated model of labor force flows and job matching, augmented with more realistic assumptions about labor force exit as a function of duration of unemployment over the business cycle.

[13] See Hobijn and Sahin (2012) for international evidence.

[14] See Daly (2012) for example.

[15] See Davis et al. (2013) on vacancy and hiring behavior at the establishment level.

3.3.1 Previous Models

In this section, I describe the Kroft et al. (2016) calibration model that I extend in the next section. Kroft and colleagues estimate a search- and matching-model that is a simplification of Mortensen and Pissarides (1994) and Shimer (2005), with slight adjustments to fit the data (for example to adjust for population growth and inconsistencies between flow data and reported durations).[16] Kroft et al. (2016) focus on workers age 25 to 54 to avoid issues concerning the aging of the baby boomers and increased school attendance. In the model there are three stocks of workers: the employed, E; the unemployed, U; and those not in the labor force, N. Between each period, which is defined as a month, agents can transition across labor force states. The probability that someone who starts in a given state finishes in a given state is given by λ^{XY}, with X the initial state and Y the final state.

For the calibration exercise, the λ's are calculated from the Current Population Survey and adjusted for biases along the lines of Shimer (2012). They are then scaled to match a peculiarity first mentioned by Elsby (2013): that many of those who enter unemployment from either employment or nonparticipation often do not report that their duration of unemployment is less than one month, which in theory it should be.[17] The probability that someone who moves from out of the labor force to unemployment reports a duration of unemployment, d, of a certain number of months is denoted $\theta_t^{NU}(d)$. Along the same lines, the probability that somebody who moves from employment to unemployment reports a duration of unemployment in months, d, is denoted $\theta_t^{EU}(d)$. These values are estimated and imposed yearly to match the stocks in the CPS. Notably, people moving into unemployment from both employment and out of the labor force the previous month reported longer durations of unemployment during and in the immediate aftermath of the recession. The transition

[16] They assume a fixed population, P, age 25-54 and adjust flows to ensure that all flows add to 100%.

[17] Those who remain unemployed in two consecutive months have another month added to their duration mechanically.

rate from unemployment to employment depends on a measure of labor force slack, total vacancies, and the duration of a worker's unemployment. The transition rate from not in the labor force to employment depends on labor force slack and vacancies. This leaves us with the following four equations that explain the movements of the labor force:

$$N_{t+1} = N_t(1 - \lambda_t^{NU} - \lambda_t^{NE}(x_t)) + E_t\lambda_t^{EN} + U_t\lambda_t^{UN} \tag{19}$$

$$U_{t+1}(0) = E_t\theta_t^{EU}(0)\lambda_t^{EU} + N_t\theta_t^{NU}(0)\lambda_t^{NU} \tag{20}$$

$$U_{t+1}(d+1) = U_t(d)(1 - \lambda_t^{UE}(x_t; d) - \lambda_t^{UN}) + E_t\theta_t^{EU}(d+1)\lambda_t^{EU} + N_t\theta_t^{NU}(d+1)\lambda_t^{NU} \tag{21}$$

$$E_{t+1} = P - U_{t+1} - N_{t+1} \tag{22}$$

where $x_t = \frac{V_t}{U+sN_t}$.

The above equations hold by construction. An important contribution of the paper is to analyze how the labor market would have evolved given the conditions of the Great Recession but under counterfactual employment flows parametrized using data from 2002 to 2007. To do this, the authors assume that the unemployed and non-participants meet job openings according to the Cobb-Douglas function

$$M(U + sN, V) = m_0(U + sN)^{\alpha}V^{1-\alpha} \tag{23}$$

where s is the weight on the search efficiency of those not in the labor force relative to the unemployed – which will be estimated from the data, V is total vacancies, and m_0 is a constant scaling parameter. In addition, the authors assume that the share of matches that go to unemployed is proportional to their share of search-intensity weighted non-employment: $\frac{U}{U+sN}$. The remaining matches go to those not in the labor force. The last assumption is that the probability that a match results in a hire for an unemployed individual depends on his/her duration of unemployment. They define this rate to be 1 for those who have been unemployed for less than one month and it decays according to a function $A(d)$, which

is estimated in three different ways, using various measures of job-finding success regressed on duration. With this estimate of $A(d)$ in hand, the authors can estimate $\hat{\lambda}_t^{UE}(x_t; d)$ and $\hat{\lambda}_t^{NE}(x_t; d)$, by choosing α and s to minimize the distance between the flows in the CPS and the following equations between 2002 and 2007, to get fitted values:

$$\hat{\lambda}_t^{UE}(x_t; d) = A(d) m_0 x_t^{1-\alpha} \tag{24}$$

$$\hat{\lambda}_t^{NE}(x_t; d) = s m_0 x_t^{1-\alpha} \tag{25}$$

where again $x_t = \frac{V_t}{U + sN_t}$. With these flows, two of the four equations for 2008 and forward become:

$$N_{t+1} = N_t(1 - \lambda_t^{NU} - \hat{\lambda}_t^{NE}(x_t; d)) + E_t \lambda_t^{EN} + U_t \lambda_t^{UN} \tag{26}$$

$$U_{t+1}(d+1) = U_t(d)(1 - \hat{\lambda}_t^{UE}(x_t; d) - \lambda_t^{UN}) + E_t \theta_t^{EU}(d+1)\lambda_t^{EU} + N_t \theta_t^{NU}(d+1)\lambda_t^{NU} \tag{27}$$

The flows out of the labor force from employment and unemployment, and the number of vacancies, are treated as exogenous forcing variables. Importantly, the unemployed are assumed to exit the labor force independently of their duration of unemployment.

Using the predicted job-finding rates for the unemployed by duration and those not in the labor force, as well as actual transition rates into non-employment and the actual path of vacancies, Kroft et al. (2016) were able to capture most of the rise in the share of long-term unemployed workers that accompanied the Great Recession. Their unemployment forecast returned to pre-recession levels rapidly, however, and they were unable to capture the observed loop in the Beveridge curve.[18] The authors note that this is because they over-

[18]This observation is based on my replication of their model. In particular, when I replicated their model I also found a rise in the share of the long-term unemployed that mirrored the observed data. When I projected the Beveridge curve using their model, there was only a slight outward shift in the curve after the Great Recession that returned to the original position by late 2012. Intuitively, this finding resulted because their meeting function generated job growth, and a consequent drop in unemployment, that was stronger than observed in the recovery because it did not allow for a lower meeting rate of the long-term unemployed and allowed the long-term unemployed to remain in the labor force for long stretches.

predict how many unemployed and not in the labor force moved into employment in each period. Additionally, their assumption that labor force exit is uncorrelated with duration breaks down over the recovery, when the long-term unemployed who stayed in the labor force longer than previously begin to withdraw rapidly. Despite these shortcomings, the model provides a useful jumping off point for investigating the loop in the Beveridge curve from a flows perspective.

3.3.2 Model Extension

In this paper, I extend the calibration model of Kroft et al. (2016) in two important respects following on the work presented in KCC. First, rather than only considering three groups of workers, I break the unemployed into two groups: U_S, the short-term unemployed workers and U_L, the long-term unemployed workers. In keeping with the rest of the text, I use 26 weeks as the cutoff duration, with those unemployed 27 weeks or longer considered long-term unemployed. The separate behavior of the long- and short-term unemployed adds two equations to the original four.

$$N_{t+1} = N_t(1 - \lambda_t^{NU} - \lambda_t^{NE}(x_t)) + E_t\lambda_t^{EN} + U_{S,t}\lambda_t^{USN} + U_{L,t}\lambda_t^{ULN} \tag{28}$$

$$U_{t+1}(0) = E_t\theta_t^{EU}(0)\lambda_t^{EU} + N_t\theta_t^{NU}(0)\lambda_t^{NU} \tag{29}$$

$$\begin{aligned} U_{t+1}(d+1) = {} & U_{S,t}(d)(1 - \lambda_t^{USE}(x_t; d) - \lambda_t^{USN}) \\ & + E_t\theta_t^{EU}(d+1)\lambda_t^{EU} + N_t\theta_t^{NU}(d+1)\lambda_t^{NU} \quad if \quad d < 6 \end{aligned} \tag{30}$$

$$U_{t+1}(7) = U_t(6)(1 - \lambda_t^{USE}(y_t; d) - \lambda_t^{USN}) + E_t\theta_t^{EU}(7)\lambda_t^{EU} + N_t\theta_t^{NU}(7)\lambda_t^{NU} \tag{31}$$

$$U_{t+1}(8) = U_t(7)(1 - \lambda_t^{ULE}(y_t; d) - \lambda_t^{ULN}) + E_t\theta_t^{EU}(8)\lambda_t^{EU} + N_t\theta_t^{NU}(8)\lambda_t^{NU} \tag{32}$$

$$E_{t+1} = P - U_{t+1} - N_{t+1} \tag{33}$$

where y_t is a measure of labor market pressure defined below, and d is duration in months. To allow the unemployed to exert differential search- and match-intensity in the model, I

then estimate a job matching function, rather than a meeting function, of the form:

$$J = (U_S = \delta U_L + \xi N)^{\alpha}(V)^{1-\alpha} \tag{34}$$

where J is the number of jobs being filled by the non-employed in a given period. . In addition to the unemployed groups defined above, N is the number of labor force nonparticipants, and V is the number of vacancies. The parameters δ and ξ reflect the potentially lower match efficiency and search-intensity of the long-term unemployed and nonparticipants. The model implies:

$$\hat{\lambda}_t^{USE}(x_t; d) = y_t^{1-\alpha} \tag{35}$$

$$\hat{\lambda}_t^{ULE}(x_t; d) = \delta y_t^{1-\alpha} \tag{36}$$

$$\hat{\lambda}_t^{NE}(x_t; d) = \xi y_t^{1-\alpha} \tag{37}$$

where $y_t = \frac{V_t}{U_S + \delta U_L + \xi N}$.

I estimate the parameters of the matching function by minimizing the distance between the system of equations for flows into employment for nonparticipants, short-term unemployed, and long-term unemployed using monthly CPS data.[19] The estimated coefficients (with bootstrapped standard errors in parentheses) are: $\delta = 0.60(.03)$, $\xi = 0.29(.01)$ and $\alpha = 0.79(.06)$.[20] This means that, on average, the long-term unemployed put 40% less pressure on the labor market and match 40% less frequently when given the chance than do the short-term unemployed. Kroft et al. (2016) implicitly assumed that $\delta = 1$ in their meeting function, although they allowed for differential job finding rates for the short- and long-term unemployed because their meeting function is combined with $A(d)$ terms to generate job

[19] To be more precise, the parameters were estimated using the labor force flows for those with more or less than 26 weeks of unemployment, and ensuring that the stocks of the number of unemployed workers implied by the transition rates match the actual CPS stocks of those with 6 and 7 months of unemployment.

[20] These are the estimates from calculations which focus on prime-aged worker alone following Kroft et al. The numbers are similar, $\delta = 0.56$, $\xi = 0.18$, and $\alpha = 0.74$ when looking at all workers age 16 or older with the estimates showing lower mobility in those groups.

matches. Their estimate of s, the match efficiency of nonparticipants, was slightly lower than my estimate of ξ, owing to their assumption of a hire for any meeting of not in labor force searchers with a vacancy. It is important to note that in this model the long-term unemployed and nonparticipants put slightly less pressure on hiring than they did in Kroft et al. (2016).

The main difference for performance of the simulation stems from the assumption imposed by Kroft et al. (2016) that the labor force withdrawal rate is the same for both the short-term and long-term unemployed.[21] While this is plausible in the immediate aftermath of a recession, the labor force withdrawal rate tends to rise as the economy recovers, especially for the long-term unemployed. Importantly, even if the average labor force withdrawal rate over the unemployed were to remain the same, the distribution matters. As the long-term unemployed withdraw, the remaining unemployed skew towards the short-term unemployed who find jobs at higher rates, further pushing down the unemployment rate. The second extension of Kroft et al (2016)'s model and simulation was to allow for differential labor force withdrawal rates by duration of unemployment over time.

3.3.3 Simulation and Updates

For the simulation assessing how well the matching function in equation (15) and the other assumptions fit the data, I follow Kroft et al. (2016) in letting the observed number of vacancies, labor force withdrawal rates, and transitions into unemployment evolve as they did from January 2008 forward. I also follow Kroft and others in assigning a duration of unemployment to those who initially transition from nonparticipant to unemployed, and those who transitioned from employed to unemployed, based on the observed distributions

[21] The discussion in Section 2 and figure 3.6, for example, show that this prediction has been largely borne out.

in that calendar year. For those who remain unemployed from one period to the next, I increment their duration of unemployment by one month.

Figure 3.9 uses equation (16) estimated over the period 2002-07 to project the Beveridge curve from 2008 to 2013 with realized ancillary flows and vacancies as was originally done in KCC. The projection seems to match the broader trends in the data reasonably well. The calibrated model predicts an outward shift in the Beveridge curve similar to what was observed, unlike in Kroft et al. (2016). In 2012 and 2013, the projection began to move back toward the original Beveridge curve. The projection initially under-predicts the 5 percentage point rise in the unemployment rate from January 2008 to October 2009 by one percentage point, and remains slightly to the left of the actual data. This under-prediction was consistent with Robert Hall and Sam Schulhofer-Wohl (2013), who found that there has been a steady downward drift in matching efficiency in the 2000s.[22] As the labor force exit rate of both the long- and short-term unemployed began to rise, the projection began to move closer to the original Beveridge curve. As of December 2013, the model predicted that the unemployment rate would be 0.8 percentage point lower than the actual rate. The model was too bullish on non-participants finding jobs. On the whole, however, the calibrated model seems to capture the broad outlines of the shift of the Beveridge curve reasonably well. The root mean square error between the actual unemployment rate and projection based on the calibrated model is 0.9 percentage point.

I next use the calibrated model to investigate the impact on unemployment of the cyclical path of labor force withdrawals by duration of unemployment. As mentioned earlier, labor force exits collapsed for the long-term unemployed during the Great Recession and in the ensuing few years, before rising in the direction of their historical average. What impact did this pattern have on the unemployment rate and shift in the Beveridge curve? Figure 3.10 presents calibrated results where I impose the labor-force withdrawal rates for the long-term

[22]Hall and Schulhofer-Wohl's analysis focuses on eight different types of job seekers to control for heterogeneity.

and short-term unemployed that occurred in 2006–07, just before the recession, each year going forward. Had labor force withdrawal rates not collapsed and instead remained at their 2006–07 levels, the figure indicates that the loop around the Beveridge curve would have been much more circumscribed and short-lived. The unemployment rate would have been under-predicted by 3 percentage points under this scenario similar to Kroft et al. (2016). The root mean squared error in this counterfactual model rises to 1.6 percentage points, more than a 60 percent worse fit than when the actual path of labor force participation by duration of unemployment is used in the model.[23]

To probe what the calibration exercise implied going forward when originally estimated, in 2014, I extend the projections through April 2018 starting with the data observed for December 2013 and using only data known in December 2013. Vacancies are treated as exogenous in the model; I assume that vacancies grow at the same rate as they had over the previous 2 years (31,000 per month). The labor force exit rates of the long- and short-term unemployed were assumed to linearly return from their December 2013 levels to their 2006 averages by 2016, an assumption that appeared consistent with the 2002–07 recovery and current trends (see figure 3.4), as were the other flows to non-employment. Figure 3.11 shows the 2014 prediction that the labor market would have almost returned to the original Beveridge curve by December 2016. This implied that the combination of rising labor force withdrawal rates and lower match efficiency for the long-term unemployed could account for a loop around the Beveridge curve. The actual performance through April 2018 shows the overall accuracy of the forecast, although the unemployment rate did not decline quite as rapidly as predicted. The mean squared error of the unemployment rate prediction was .50 percentage point through 2016. In addition to the unemployment rate, the model undershot the vacancy rate. Vacancies remain at record levels and growing into 2018, something that seemed unlikely at the end of 2013. Figure 3.11 allows for a better optical comparison by

[23]See also Figure 21 from the paper which decomposed the effects of this rising exit rate and showed it had explained 30 percent of the unemployment rate decline through the end of 2013.

imposing the predicted unemployment rates on the actual path of vacancies. The results are much more accurate than consensus at the time when many believed that unemployment was at its natural rates.

As was noted, my preferred forecasts were made for those aged 25-54 to abstract away from changes in demographics, especially the retirement of the baby boomers, and increased schooling. This prevents me from making an apples-to-apples comparison of errors with other forecasters, who generally only predict the overall unemployment rate. To rectify this, I also ran the simulation for all those 16 years or older. The model does a much worse job of fitting post-2007 dynamics overall, especially at predicting the increase in out of the labor force. Table 3.1, however, shows that the model using data through March 2014 still produced an overall unemployment rate forecast that had a lower absolute error than the Federal Reserve Forecast or the Survey of Professional Forecasters. The projection does more than 15% better than both the median forecast from the Survey of Professional Forecasters and the Central Tendency estimates from the Federal Reserve Board Members and the Federal Reserve Bank Presidents.[24] The forecast of a rapid decline in the unemployment rate, was well below the mainstream view. Although the model was too bullish on nonparticipants finding jobs and undershot the unemployment rate overall, the flows approach with an emphasis on the long-term unemployed did better than other available forecasts. Moreover, the fact that a similar pattern was observed in the past recovery underscores that the labor force withdrawal of the long-term unemployed is historically an important (but unfortunate) mechanism by which the labor market returns loops around the Beveridge Curve.

[24]For forecasts more than four quarters in the future, the Survey only reports annual values, but those forecasts were too high by more than a percentage point.

3.4 Summary and Conclusion

From its previous range of 10 to 20 percent of total unemployment, long-term unemployment reached 45 percent of total unemployment in the U.S. Concurrently, many of the stable relationships in the economy involving the unemployment rate appeared to have broken down. The Beveridge curve, which traces the relationship between unemployment and vacancies, seemed to have shifted out substantially. This led many to suggest that this time was different, that the job market was permanently changed. Stories were told by some about how automation created less need for workers, especially the long-term unemployed who were assumed to be the relatively less productive workers. Other commenters predicted that workers were just retooling and otherwise waiting out the recession and that there would be an influx of workers to the labor force if the economy continued to strengthen. Neither of these explanations for the initial shift in the Beveridge curve has withstood the test of time.

This paper demonstrates that it was the cyclical behavior of the record number of long-term unemployed during and after the recession that explained the initial movements away from the Beveridge curve and led to the completion of a counterclockwise loop as predicted in Blanchard and Diamond (1989). To quantify the magnitude of the labor market effects of the long-term unemployed throughout the business cycle precisely, I extend the calibrated model originally estimated in Krueger, Cramer, and Cho (2014) and focus on the channels through which duration dependence could be a factor. Incorporating the cyclical patterns in the labor force flows of the long-term unemployed (lower labor force withdrawal rates during recessions and higher labor force withdrawal rates as recoveries continue) improved the accuracy of the model's forecast during the recession by 60 percent and resulted in forecasts which were 15 percent more accurate than the Federal Reserve and the Survey of Professional Forecasters during the recovery. This paper also documents the continued poor performance of the previously long-term unemployed in the face of a lasting recovery,

who were still more than twice as likely to exit the labor force as to find employment in 2018. This underscores the need for policymakers to intervene to prevent the human and economic tragedy that is long-term joblessness before it takes root fully. Unfortunately, little progress has been made and these workers continue to be at the margins of the labor market, struggling well into the ninth year of the recovery.

References

Abowd, John M., and Arnold Zellner. 1985. “Estimating Gross Labor-Force Flows.” Journal of Business & Economic Statistics 3, no. 3: 254–83.

Abraham, Katharine, and Robert Shimer. 2002. “Changes in Unemployment Duration and Labor Force Attachment.” In The Roaring Nineties, edited by Alan Krueger and Robert Solow. New York: Russell Sage Foundation.

Barnichon, Regis, and Andrew Figura. 2013. “Declining Labor Force Attachment and Downward Trends in Unemployment and Participation.” Working Paper no. 2013-88, Finance and Economics Discussion Series. Washington: Board of Governors of the Federal Reserve System.

Blanchard, Olivier Jean, and Peter Diamond. 1989. “The Beveridge curve.” Brookings Papers on Economic Activity no. 1: 1 – 76.

_____. 1994. “Ranking, Unemployment duration, and Wages.” The Review of Economic Studies, 61(3), pp.417-434.

Blanchard, Olivier Jean, and Lawrence H. Summers. 1987. “Hysteresis in Unemployment.” European Economic Review 31, no. 1: 288–95.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2012. "Recruiting intensity during and after the Great Recession: National and industry evidence." American Economic Review 102, no. 3 (2012): 584-88.

_____. 2013. "The establishment-level behavior of vacancies and hiring." The Quarterly Journal of Economics 128, no. 2 (2013): 581-622.

DeLong, J. Bradford, and Lawrence H. Summers. 2012. "Fiscal Policy in a Depressed Economy." Brookings Papers on Economic Activity, Spring: 233–97.

Elsby, Michael W.L., Bart Hobijn, Aysegül Sahin, and Robert G. Valletta. 2011. "The Labor Market in the Great Recession—An Update to September 2011." Brookings Papers on Economic Activity, Fall: 353–84.

_____. 2013. "On the Importance of the Participation Margin for Labor Market Fluctuations." Working Paper no. 2013-05. San Francisco: Federal Reserve Bank of San Francisco.

Elsby, Michael W.L., Ryan Michaels and Gary Solon. 2009. "The Ins and Outs of Cyclical Unemployment." American Economic Journal: Macroeconomics no. 1: 84–110.

Farber, Henry. 2015. "Job Loss in the Great Recession and its Aftermath: U.S. Evidence from the Displaced Workers Survey? Evidence from Recent Cycles in the US Labor Market." Working Paper no. 589. Princeton, N.J..: Princeton Industrial Relations Section.

Farber, Henry S., and Robert G. Valletta. 2015. "Do extended unemployment benefits lengthen unemployment spells? Evidence from recent cycles in the US labor market." Journal of Human Resources 50, no. 4 (2015): 873-909.

Forslund, Anders, and Alan B. Krueger. "An evaluation of the Swedish active labor market policy: new and received wisdom." In The Welfare State in Transition: Reforming the Swedish Model, pp. 267-298. University of Chicago Press, 1997.

_____. 2010 "Did Active Labor Market Policies Help Sweden Rebound from the Depression of the Early 1990s?." In Reforming the welfare state: Recovery and beyond in Sweden, pp. 159-187. University of Chicago Press.

Ghayad, Rand. 2013. “The Jobless Trap” Working paper (http://media.wix.com/ugd/576e9a_f6cf3b6661e44621ad26547112f66691.pdf).

Gordon, Robert J. 2013. The Phillips curve is alive and well: Inflation and the NAIRU during the slow recovery. No. w19390. National Bureau of Economic Research.

Hall, Robert E. and Sam Schulhofer-Wohl. 2013. “Measuring Matching Efficiency with Heterogeneous Jobseekers.” Mimeo. Stanford University.

Hobijn, Bart, and Aysegül Sahin. 2012. “Beveridge curve Shifts across Countries Since the Great Recession.” Paper presented at the International Monetary Fund’s 13th Jacques Polak Annual Research Conference, Washington, DC, November 8–9.

IMF (2010), “Unemployment Dynamics in Recessions and Recoveries,” World Economic Outlook, Chapter 3, April 2010.

Katz, Lawrence F. 1986. “Layoffs, Recall, and the Duration of Unemployment.” Working Paper no. 1825. Cambridge, Mass.: National Bureau of Economic Research.

Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo. 2013. “Duration Dependence and Labor Market Conditions.” Quarterly Journal of Economics 128, no. 3: 1123–167.

Kroft, Kory, Fabian Lange, Matthew J. Notowidigdo, and Lawrence F. Katz. "Long-term unemployment and the Great Recession: the role of composition, duration dependence, and nonparticipation." Journal of Labor Economics 34, no. S1 (2016): S7-S54.

Kroft, Kory, Fabian Lange, Matthew J. Notowidigdo, and Matthew Tudball. “Long time out: unemployment and joblessness in Canada and the United States.” Working paper. http://faculty. wcas. northwestern. edu/noto/research/KLNT_ June_2017. pdf, 2017.

Krueger, Alan B. "Where Have All the Workers Gone?: An Inquiry into the Decline of the US Labor Force Participation Rate." Brookings papers on economic activity 2017, no. 2 (2017): 1-87.

Krueger, Alan B., and Andreas Mueller. 2011. “Job Search, Emotional Well-Being, and Job Finding in a Period of Mass Unemployment: Evidence from High-Frequency Longitudinal Data.” Brookings Papers on Economic Activity, Spring: 1–57.

Krueger, Alan B., Judd Cramer, and David Cho. "Are the long-term unemployed on the margins of the labor market?." Brookings papers on economic activity 2014, no. 1 (2014): 229-299.

Mortensen, Dale T., and Christopher A. Pissarides. 1994. “Job Creation and Job Destruction in the Theory of Unemployment.” Review of Economic Studies 61, no. 3: 397 – 415.

Mueller, Andreas. 2012. “Separations, Sorting and Cyclical Unemployment.” IZA Discussion Paper no. 6849.

Mueller, Andreas, Jesse Rothstein, Till von Wachter. 2013. “Unemployment Insurance and Disability Insurance in the Great Recession.” Working Paper no. 19672. Cambridge, Mass.: National Bureau of Economic Research.

Nekarda, Christopher J. 2009. “A Longitudinal Analysis of the Current Population Survey: Assessing the Cyclical Bias of Geographic Mobility.” Working Paper (http://chrisnekarda.com/papers/geographic-mobility.pdf) (with data updated through December 2013).

Poterba, James M. and Lawrence H. Summers. 1986. “Reporting Errors and Labor Market Dynamics.” Econometrica Society 54(6): 1319-1338.

Shan, Hui and Stehn, Jari. “Young vs. Old as Drivers of Labor Force Participation.” Goldman Sachs Global Macro Research. 2014.

Shimer, Robert. 2012. “Reassessing the Ins and Outs of Unemployment.” Review of Economic Dynamics 15, NO. 2: 127–48.

_____. 2005. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies.” American Economic Review 95, NO. 1: 25–49.

Stevens, Ann Huff. 1997. “Persistent Effects of Job Displacement: The Importance of Multiple Job Losses.” Journal of Labor Economics 15, no. 1: 165 – 88.

Stock, James H. 2011. “Discussion of Ball and Mazumder.” Brookings Papers on Economic Activity (Spring): 387 – 402.

Valletta, Robert G. 2011. “Rising Unemployment Duration in the United States: Composition or Behavior?” Mimeo. San Francisco: Federal Reserve Bank of San Francisco.

Wanberg, Connie R., Jing Zhu, Ruth Kanfer, and Zhen Zhang. 2012. “After the Pink Slip: Applying Dynamic Motivation Frameworks to the Job Search Experience.” Academy of Management Journal 55: 261-264.

Yellen, Janet L. "Labor market dynamics and monetary policy." In Remarks given at the Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming, August, vol. 22. 2014.

Tables

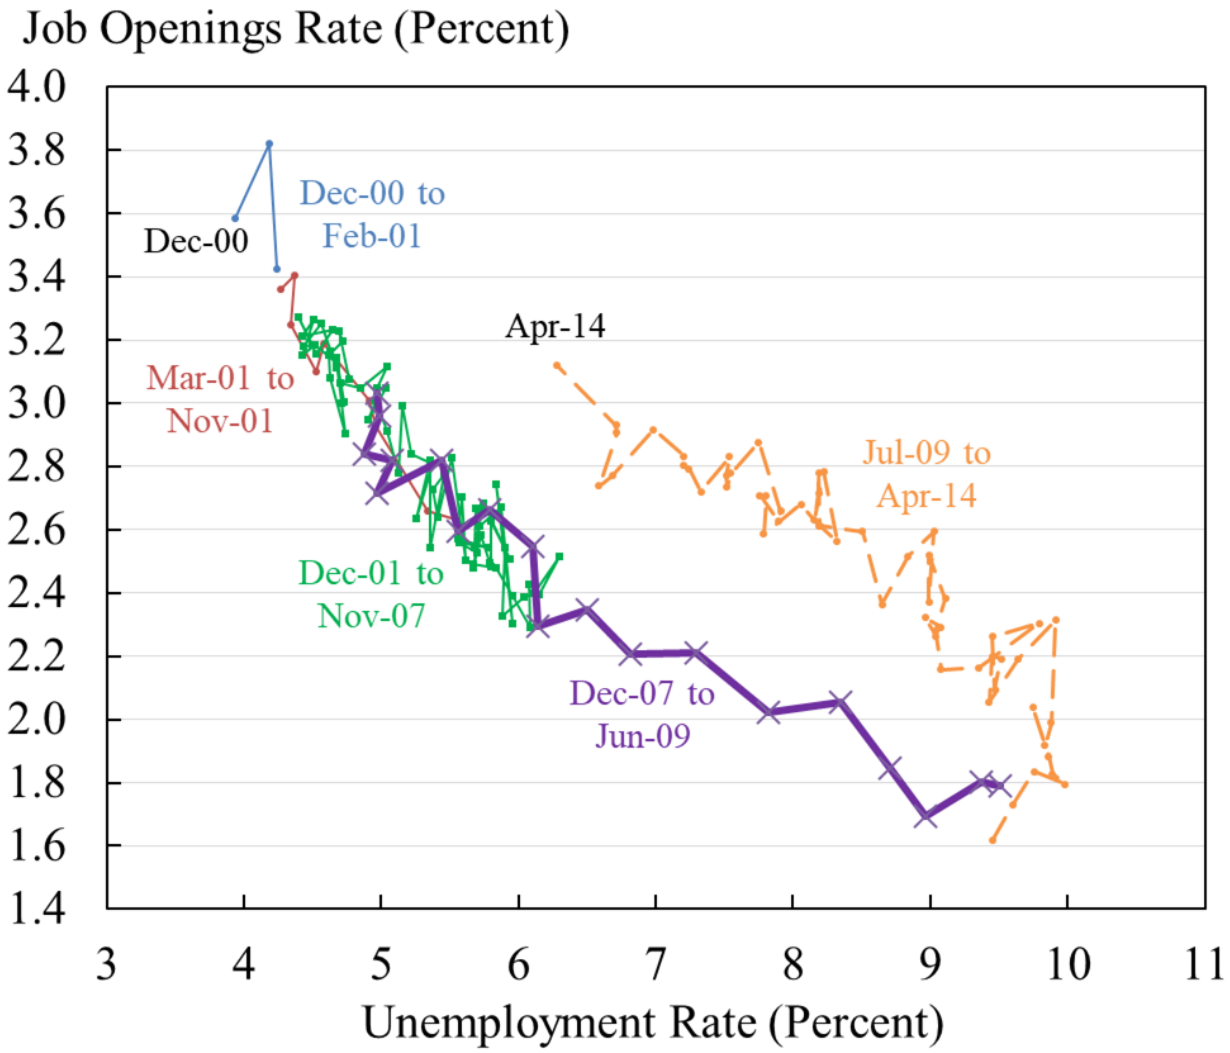
Table 3.1: Forecasts of the Unemployment Rate and Actual Rate For All Workers 16+, 2014-2017

	Year			
	2014	**2015**	**2016**	**2017**
Actual[1]	**6.2**	**5.3**	**4.9**	**4.4**
Survey of Professional Forecasters (2014)[2]	**6.4**	**5.9**	**5.6**	**5.5**
Federal Reserve (2014)[3]	**6.1-6.3**	**5.6-5.9**	**5.2-5.6**	**5.2-5.6**
Current Paper[4]	**6.0**	**4.9**	**4.3**	**3.9**

[1] Annual Average from Bureau of Labor Statistics
[2] Third Quarter 2014, Annual Projection, available at: https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/2014/survq314
[3] March 2014, Economic Projections of Central Tendency, 2017 Value corresponds to "longer run" forecast
[4] Annual Average, alternative forecast for all workers 16 years and older

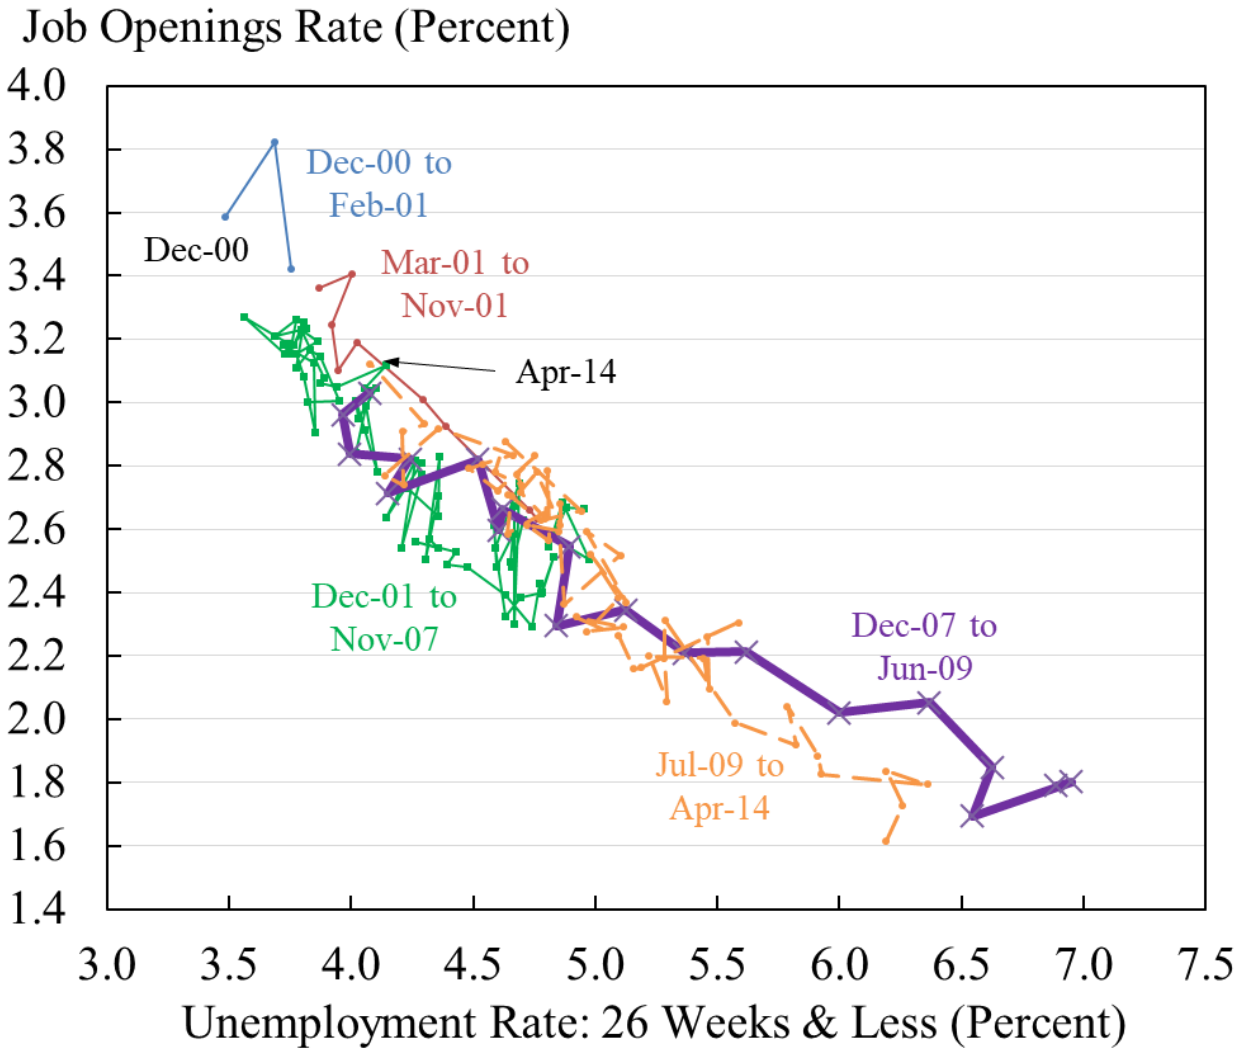
Figures

Figure 3.1: (a) Job Openings vs. Unemployment Rate, Dec-00 to Apr-14



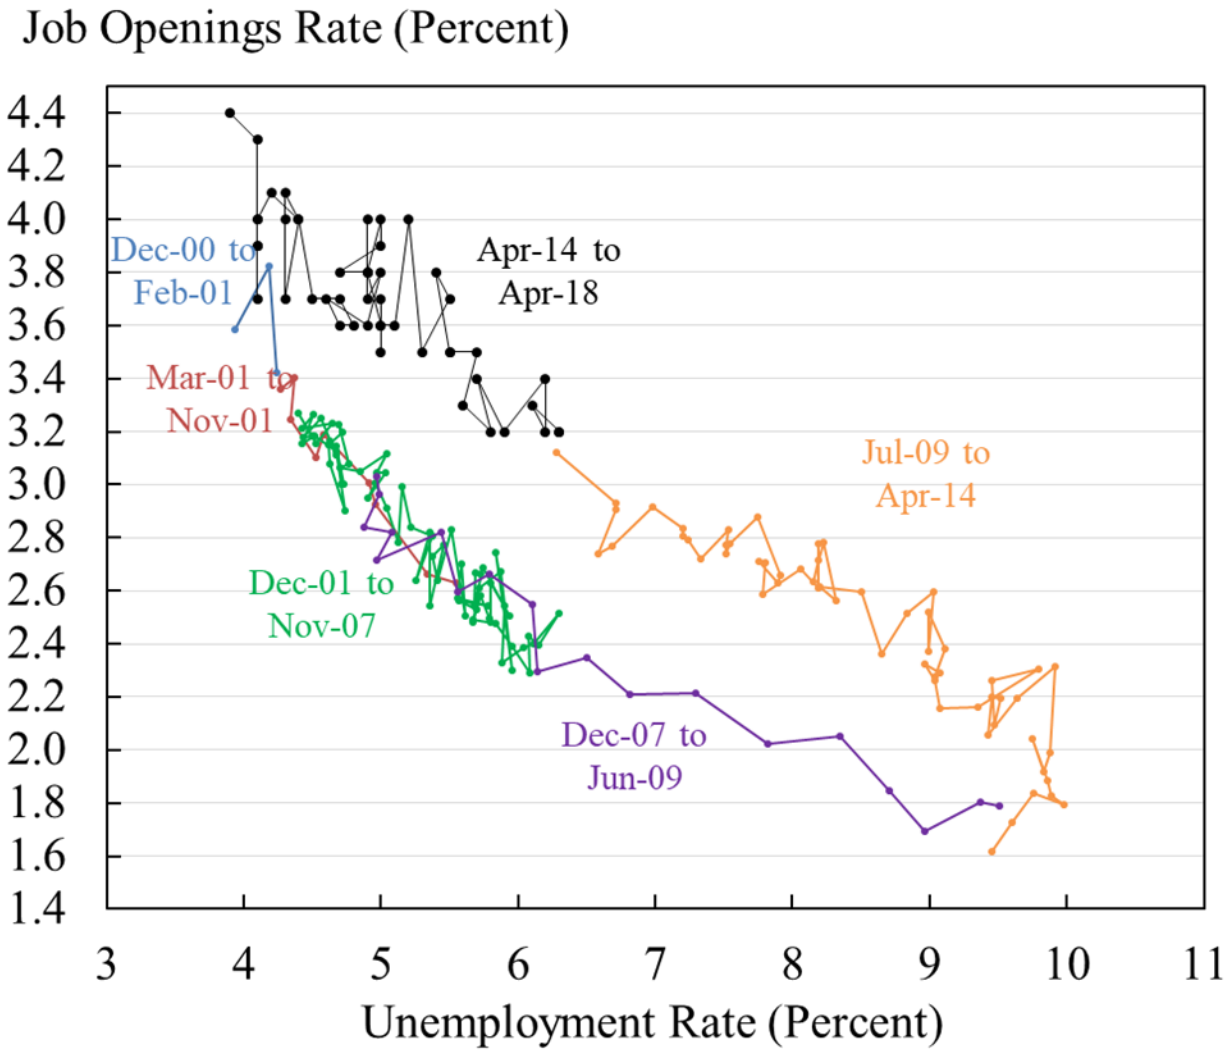
Note: Job openings rate is defined as job openings as a percentage of the sum of job openings and total nonfarm payroll employment.
Source: Bureau of Labor Statistics (Job Openings and Labor Turnover Survey and Current Population Survey).

Figure 3.1: (b) Job Openings vs. Short-Term Unemployment Rate, Dec-00 to Apr-14



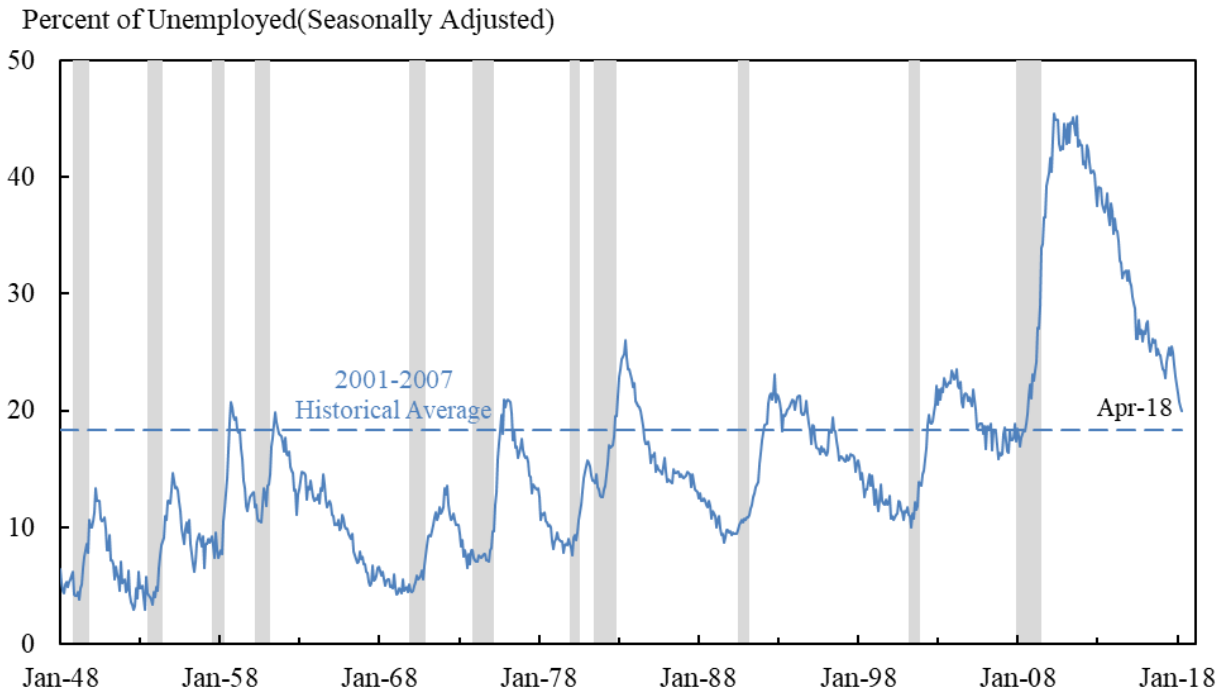
Note: Job openings rate is defined as job openings as a percentage of the sum of job openings and total nonfarm payroll employment.
Source: Bureau of Labor Statistics (Job Openings and Labor Turnover Survey and Current Population Survey).

Figure 3.2: Job Openings vs. Unemployment Rate, Dec-00 to Apr-18



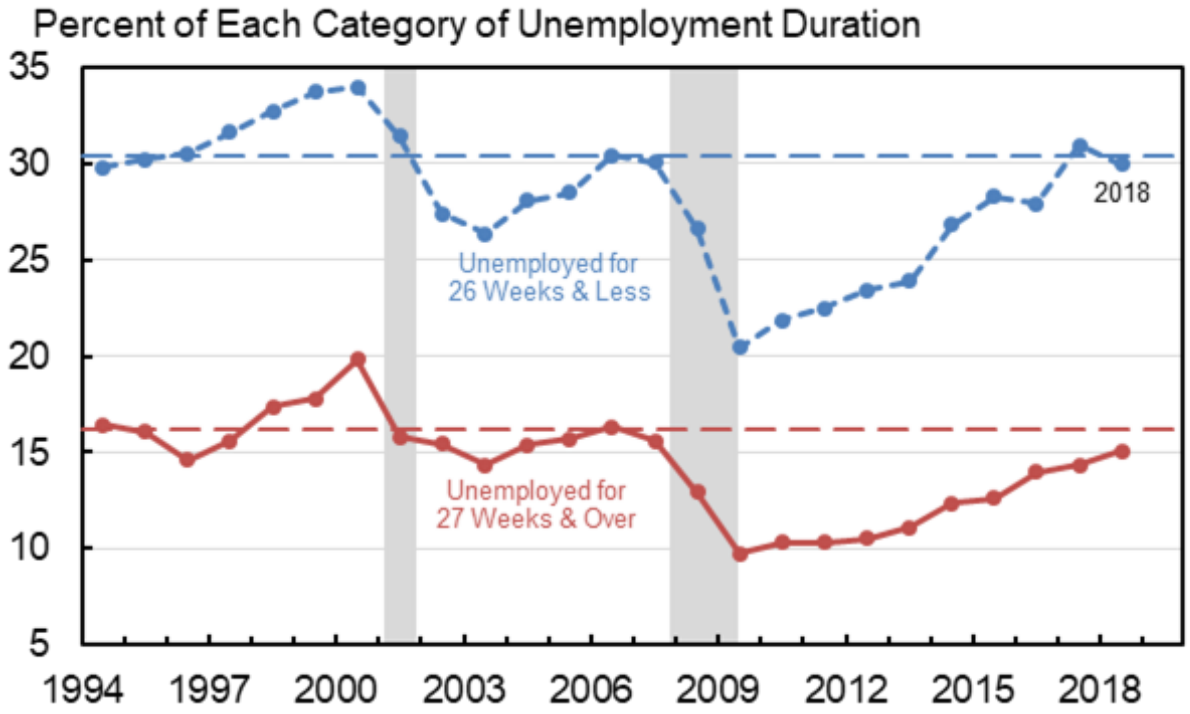
Note: Job openings rate is defined as job openings as a percentage of the sum of job openings and total nonfarm payroll employment.
Source: Bureau of Labor Statistics (Job Openings and Labor Turnover Survey and Current Population Survey).

Figure 3.3: Of Total Unemployed, Share Unemployed over 26 Weeks



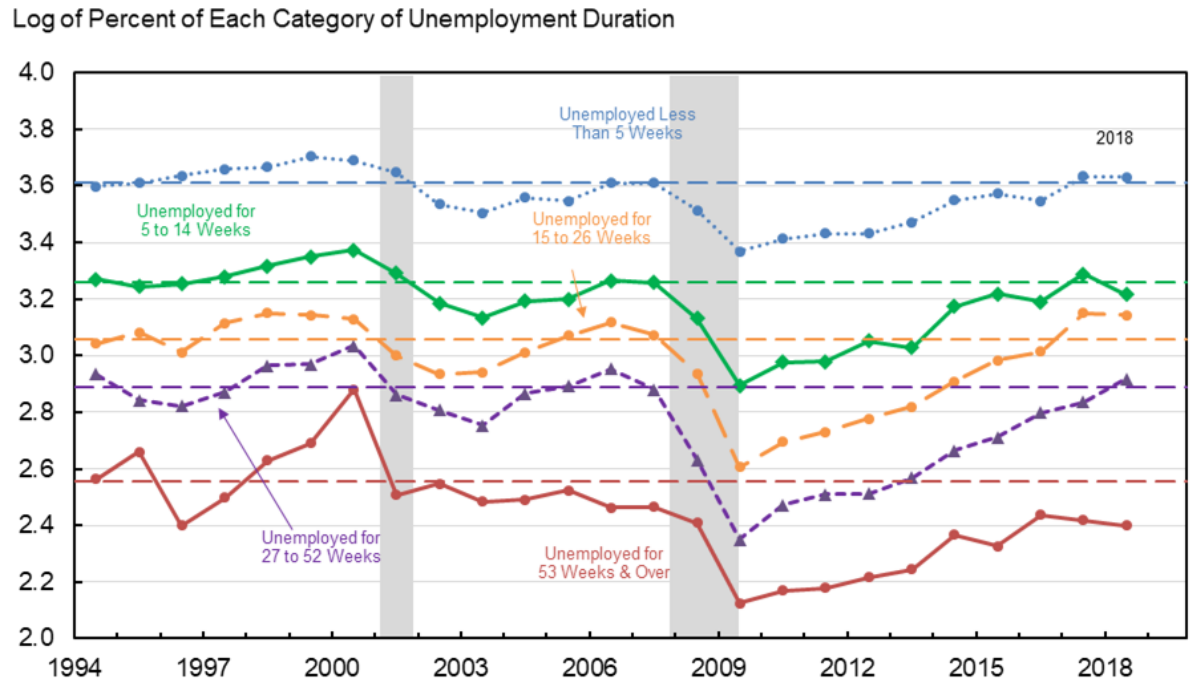
Note: Shading denotes recession.
Source: Bureau of Labor Statistics (Current Population Survey); National Bureau of Economic Research.

Figure 3.4: Monthly Probability of Transitioning from Unemployment to Employment by Duration of Unemployment



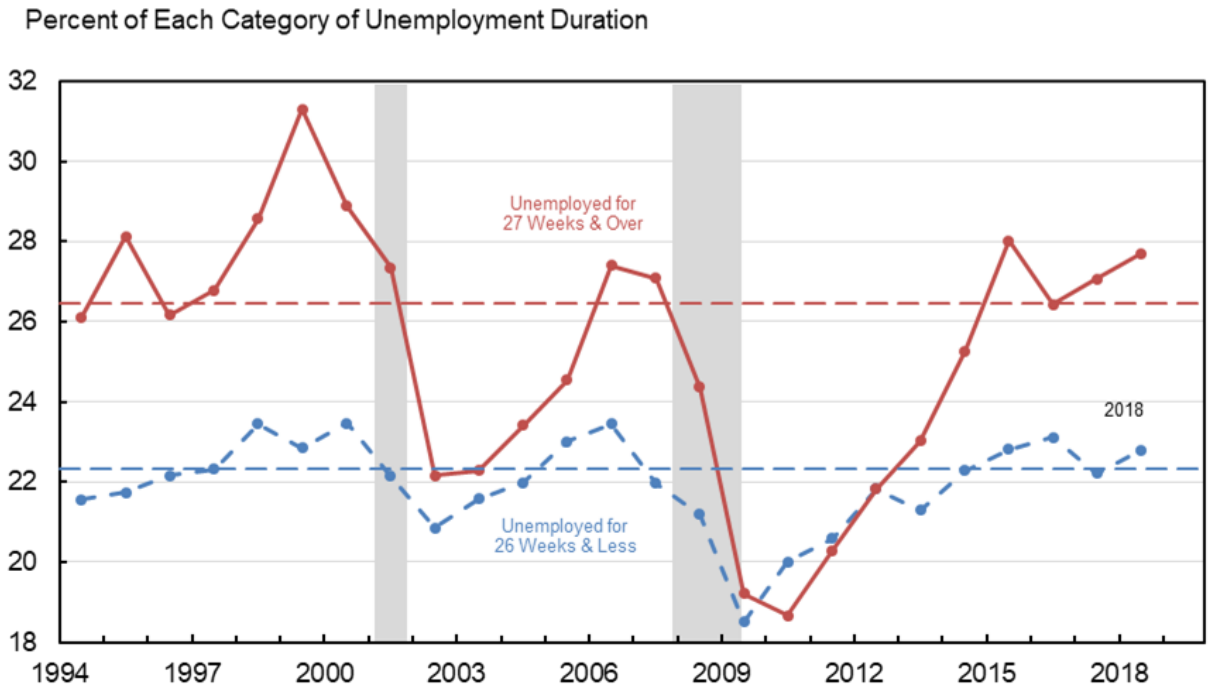
Note: Dashed lines represent 1994-2007 averages. Shading denotes recession.
Source: Bureau of Labor Statistics (Current Population Survey); National Bureau of Economic Research.

Figure 3.5: Monthly Probability of Transitioning from Unemployment to Employment by Duration of Unemployment



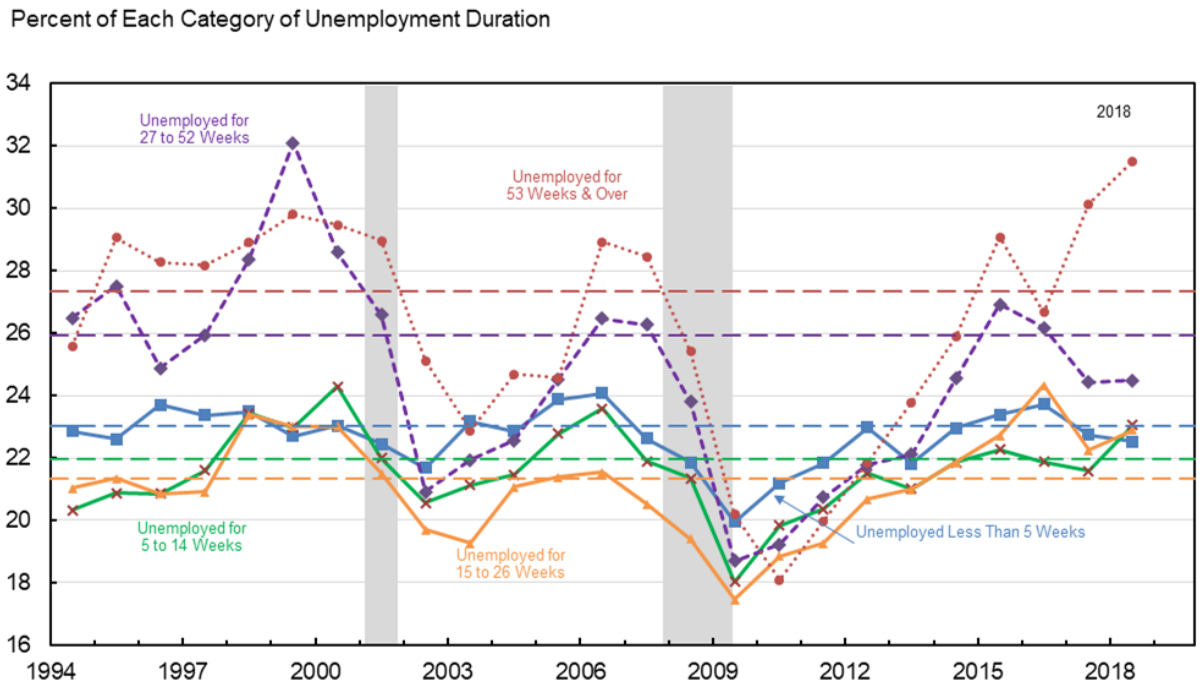
Note: Dashed lines represent 1994-2007 averages. Shading denotes recession.
Source: Bureau of Labor Statistics (Current Population Survey); National Bureau of Economic Research.

Figure 3.6: Monthly Probability of Transitioning from Unemployment to Out of the Labor Force by Duration of Unemployment



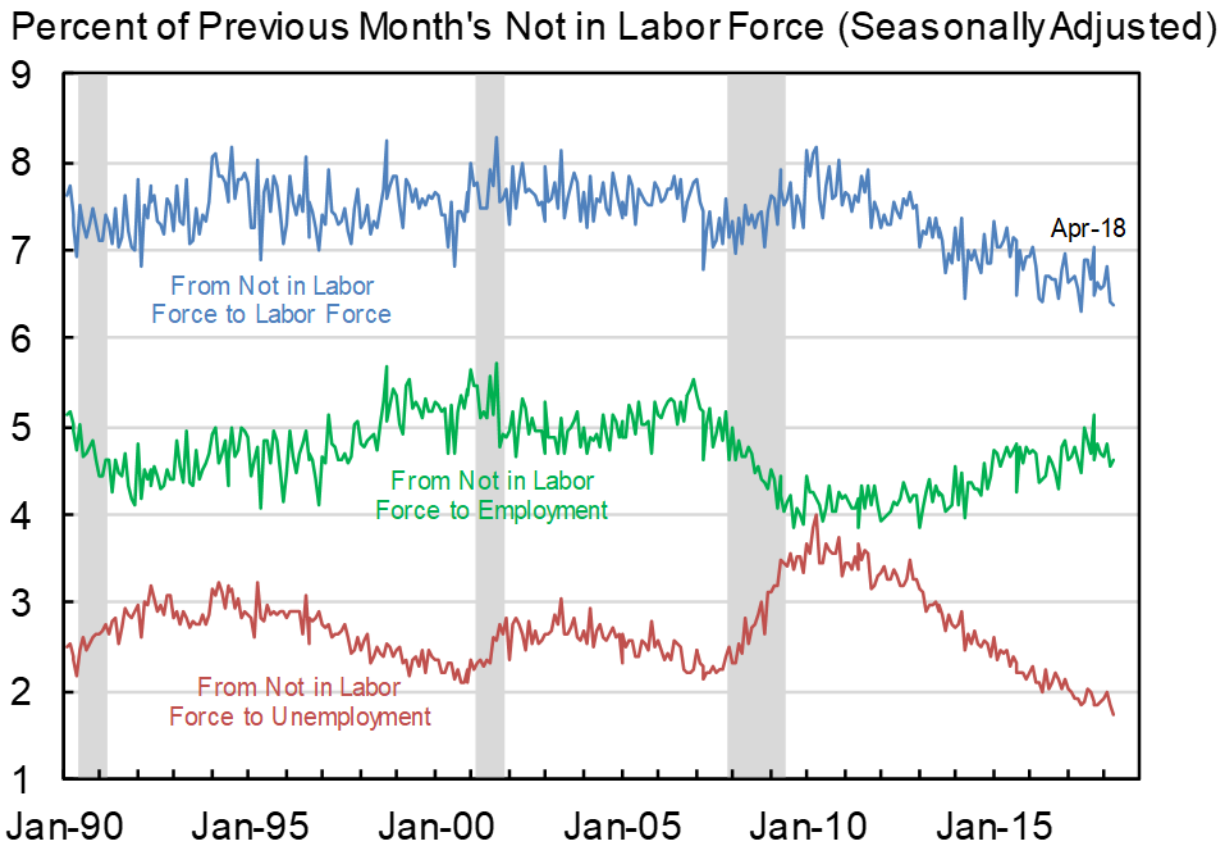
Note: Dashed lines represent 1994-2007 averages. Shading denotes recession.
Source: Bureau of Labor Statistics (Current Population Survey); National Bureau of Economic Research.

Figure 3.7: Monthly Probability of Transitioning from Unemployment to Out of the Labor Force by Duration of Unemployment



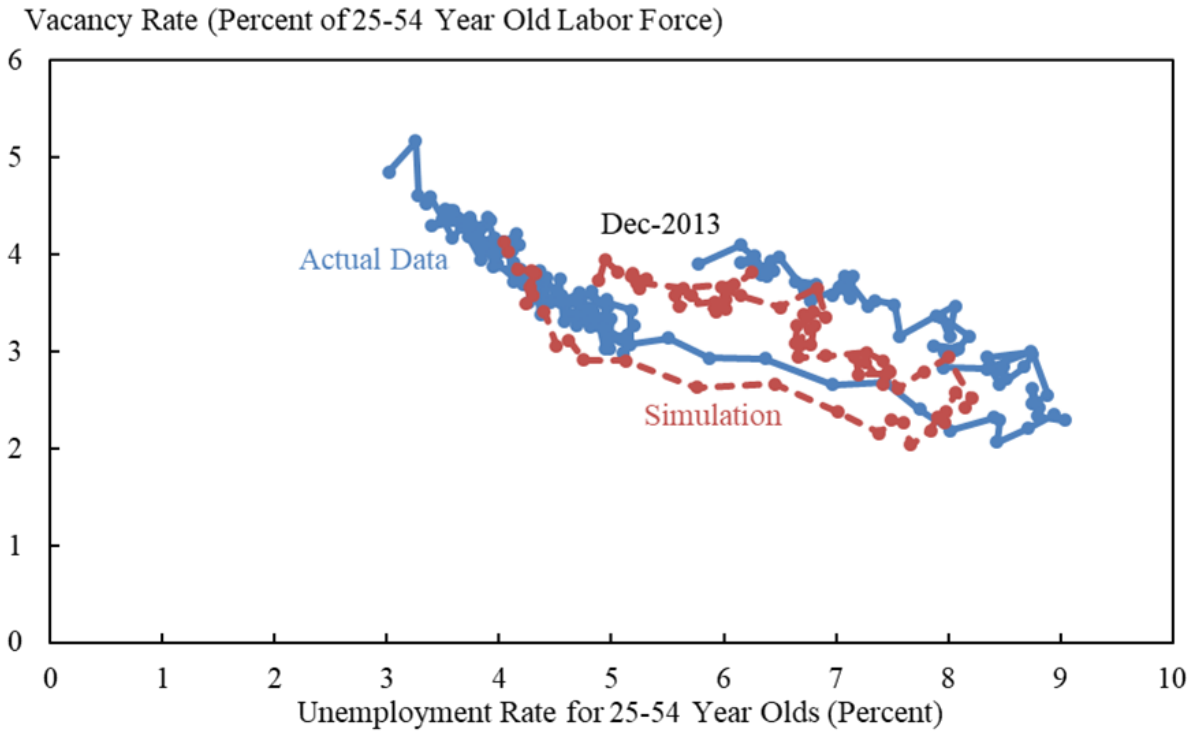
Note: Dashed lines represent 1994-2007 averages. Shading denotes recession.
Source: Bureau of Labor Statistics (Current Population Survey); National Bureau of Economic Research.

Figure 3.8: Transition Rate from Not in Labor Force



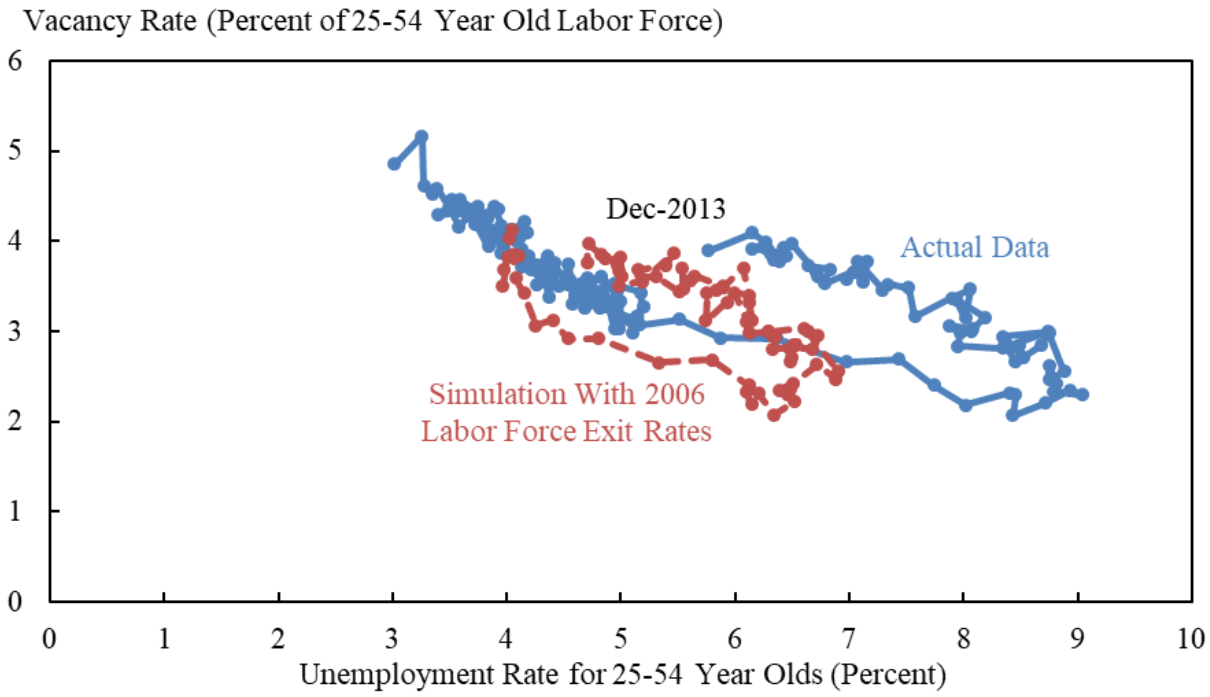
Note: Shading denotes recession.
Source: Bureau of Labor Statistics (Current Population Survey); National Bureau of Economic Research.

Figure 3.9: Calibrated Beveridge Curve, 2008-2013



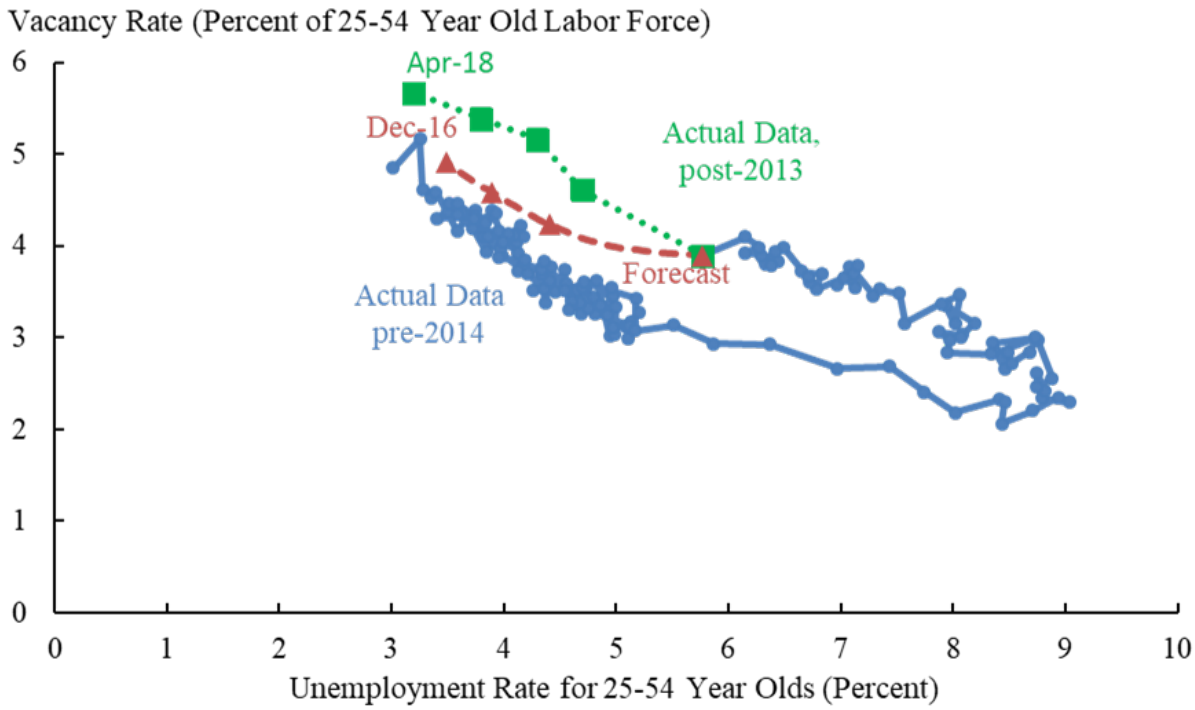
Source: Bureau of Labor Statistics (Current Population Survey and Job Openings and Labor Turnover Survey); author's calculations.

Figure 3.10: Calibrated Beveridge Curve with 2006 to 2007 Average Labor Force Exit Rates



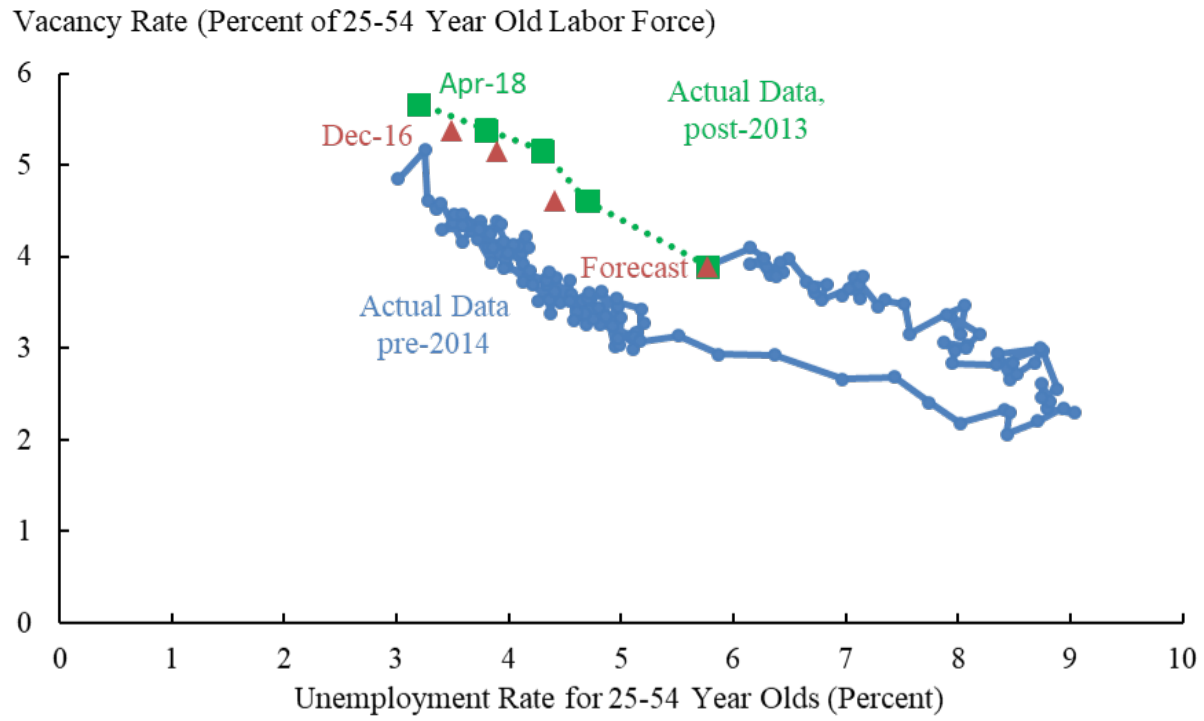
Source: Bureau of Labor Statistics (Current Population Survey and Job Openings and Labor Turnover Survey); author's calculations.

Figure 3.11: KCC Beveridge Curve Projection and Actual Data



Source: Bureau of Labor Statistics (Current Population Survey and Job Openings and Labor Turnover Survey); author's calculations.

Figure 3.12: KCC Beveridge Curve Projection and Realized Vacancy Rate and Actual Data



Source: Bureau of Labor Statistics (Current Population Survey and Job Openings and Labor Turnover Survey); author's calculations.